Bilkent Üniversitesi
Mühendislik Fakültesi
Endüstri Mühendisliği Bölümü

Üniversite-Sanayi İşbirliği Öğrenci Projeleri 2026

Derleyenler

Prof. Dr. Savaş Dayanık, Doç. Dr. Emre Nadar,
Dr. Emre Uzun
Endüstri Mühendisliği Bölümü
Sistem Tasarımı Ders Koordinatörleri

Yeşim Gülseren
Üniversite-Sanayi İşbirliği Öğrenci Projeleri
Koordinatörü

BASKI: Meteksan Matbaacılık, Mayıs 2026
YAYINCI SERTİFİKA NO: 51344
MATBAA SERTİFİKA NO: 46519

Bandrol Uygulamasına İlişkin Usul ve Esaslar Hakkında Yönetmeliğin 5. maddesinin 2. fıkrası çerçevesinde bandrol taşıması zorunlu değildir.

According to the second paragraph of Article 5 of the Regulation on Banderole Application Procedures and Principles, the use of banderole for this publication is not compulsory.

Para ile satılamaz.
For free distribution only.

Bilkent University
Faculty of Engineering
Department of Industrial Engineering

University-Industry Collaboration Student Projects 2026

Editors

Prof. Dr. Savaş Dayanık, Assoc. Prof. Emre Nadar,
Dr. Emre Uzun
Department of Industrial Engineering
Systems Design Course Coordinators

Yeşim Gülseren
University-Industry Collaboration Student Projects
Coordinator

Düzenleme Kurulu

Organizing Committee

Aslı Baytar
Prof. Dr. Savaş Dayanık
Semra Selin Eraslan
Fatih Selim Erdem
Doç. Dr. Emre Nadar
Yeşim Gülseren
Ece Kuşdemir
Dr. Emre Uzun

Değerlendirme Kurulu

Review Committee

Bilkent Üniversitesi Endüstri Mühendisliği Öğretim Üyeleri

Bilkent University Industrial Engineering Faculty Members

Prof. Dr. M. Selim Aktürk
Doç. Dr. Çağın Ararat
Prof. Dr. Savaş Dayanık
Prof. Dr. Nesim K. Erkip
Prof. Dr. Ülkü Gürler
Doç. Dr. Özlem Çavuş İyigün
Dr. Öğr. Üyesi Ayşenur Karagöz
Prof. Dr. Oya Ekin Karaşan
Doç. Dr. Ayşe Selin Kocaman
Doç. Dr. Emre Nadar
Dr. Öğr. Üyesi Gizem Özbaygın
Prof. Dr. Mustafa Çelebi Pınar
Prof. Dr. Alper Şen
Doç. Dr. Firdevs Ulus
Dr. Emre Uzun
Prof. Dr. Bahar Yetiş Kara

Doç. Dr., *Assoc. Prof.*; Dr. Öğr. Üyesi, *Asst. Prof.*

Önsöz

Bu kitap, 2025-2026 öğretim yılında Bilkent Üniversitesi Endüstri Mühendisliği Bölümü tarafından gerçekleştirilen *Üniversite-Sanayi İşbirliği Bitirme Projeleri* özetlerini kapsamaktadır. Programımız 32 yıl önce sistem tasarımı derslerinin sanayi projelerine dönüştürülmesi ile başlamıştır. Bu süre içerisinde, farklı sektör ve ölçekte faaliyet gösteren 145 iş, sanayi, ve kâr amacı gütmeyen kuruluşla toplam 599 proje gerçekleştirilmiştir.

Endüstri Mühendisliği Bölümü dördüncü sınıf öğrencilerinden oluşan proje ekipleri, akademik ve iş dünyasından danışmanların gözetiminde kurumların gündemine girmiş olan ve çözüm bekleyen gerçek problemlerini çözmektedirler. Yapılan projeler sonucunda ortaya çıkan ürün, yöntem veya hizmet, ilgili kurumlara önemli yarar ve katma değer sağlamaktadır.

Endüstri Mühendisliği Proje Fuarı ve Yarışması, 2003 yılında yapılan projelerin ilgili tüm firma, kuruluş ve üniversitelerle paylaşılması, iş dünyasının seçkin kuruluşlarının birbirleriyle ve üniversite ile olan etkileşiminin artırılması ve öğrencilerimizin iş hayatına daha donanımlı hazırlanmasını sağlamak amacıyla başlatılmıştır. Her yıl sistematik ve etkin bir şekilde yapılan bu çalışmaların daha kalıcı ve yaygın olarak paylaşılması amacıyla da "Endüstri Projeleri" kitabı serisi hazırlanmış ve bu dönemde gerçekleştirilen projeler gizlilik ilkesine bağlı kalınarak özet halinde sizlere sunulmuştur.

Kitapta yer alan proje özetlerinin doğru ve okunaklı olması için desteklerini esirgemeyen *Değerlendirme Kurulu*'muza, fuar ve yarışma jürimizde görev alan Dr. Alptekin Demiray (ORS), Dr. Bora Dilik (Nevzat Ecza), Mehmet Mustafa Tanrıkulu (Memuta Süt), Hande Dilek Yiğitel (Aselsan) ve Prof. Dr. Ülkü Gürler'e (Bilkent Üniversitesi) teşekkür ederiz.

Prof. Dr. Savaş Dayanık
Doç. Dr. Emre Nadar
Dr. Emre Uzun

Bilkent Üniversitesi Endüstri Mühendisliği Bölümü
Sistem Tasarımı Dersi Koordinatörleri

Preface

This booklet contains 2025-2026 academic year *University-Industry Collaboration Student Project* summaries done by the senior students of the Industrial Engineering Department at Bilkent University in collaboration with industrial companies, businesses, and non-profit organizations. This program started when senior design courses were reorganized as industrial projects 32 years ago. Since then, 599 projects have been completed, with 145 companies operating in various sectors.

Senior student groups of the Industrial Engineering Department solve companies' real problems under the guidance of academic and industrial advisors. The project outcomes provide companies with many operational benefits and add value to their services and products.

Since 2003 *Industrial Engineering Project Fair and Competition* has been held to disseminate the project outcomes to firms and universities, boost the synergy, encourage collaboration between industry and university, and help senior students get better equipped before they take full industrial positions. Every year the project summaries are edited in a project booklet with care given not to disclose firm-specific sensitive information and shared with the community to spread the word and impact of projects.

We thank the *Review Committee* for their efforts that improved the correctness and readability of project summaries in the book. We also thank Dr. Alptekin Demiray (ORS), Dr. Bora Dilik (Nevzat Ecza), Mehmet Mustafa Tanrıkulu (Memuta Süt), Hande Dilek Yiğitel (Aselsan), and Prof. Dr. Ülkü Gürler (Bilkent University) for serving on the project competition jury this year.

Prof. Dr. Savaş Dayanık
Assoc. Prof. Emre Nadar
Dr. Emre Uzun

Bilkent University Industrial Engineering Department
Systems Design Course Coordinators

İçindekiler
Table of Contents

Bugüne kadar öğrenci projelerimize destek veren kuruluşlar

Companies participated in the student projects so far

Düzenleme kurulu, 2025-2026 programına değerli katkıları için aşağıda adı geçen Bilkent Üniversitesi mensuplarına teşekkür eder.

The organizing committee thanks Bilkent University members named below for their invaluable help to run 2025-2026 program.

Endüstri Mühendisliği Öğretim Üyeleri

Industrial Engineering Faculty Members

Prof. Dr. M. Selim Aktürk
Doç. Dr. Çağın Ararat
Prof. Dr. Savaş Dayanık
Prof. Dr. Nesim K. Erkip
Prof. Dr. Ülkü Gürler
Doç. Dr. Özlem Çavuş İyigün
Dr. Öğr. Üyesi Ayşenur Karagöz
Prof. Dr. Oya Ekin Karaşan
Doç. Dr. Ayşe Selin Kocaman
Doç. Dr. Emre Nadar
Dr. Öğr. Üyesi Gizem Özbaygın
Prof. Dr. Mustafa Çelebi Pınar
Prof. Dr. Alper Şen
Doç. Dr. Firdevs Ulus
Dr. Emre Uzun
Prof. Dr. Bahar Yetiş Kara

Üniversite Sanayi İşbirliği Öğrenci Projeleri Koordinatörü

University & Industry Cooperation Student Projects Coordinator

Yeşim Gülseren

Endüstri Mühendisliği Bölümü İdari Asistanı

Industrial Engineering Dept. Administrative Assistant

Vegün Ekmekçi

Araştırma, Planlama ve Koordinasyon Birimi

Research, Planning and Coordination Unit

Sinan Atalar
Nuray Güvenir
Handan Keskin

Mali İşler Müdürlüğü

Financial Affairs Office

Ramazan Atmaca
Ömer Faruk Can
Arda Kaya
Serdar Ser

Düzenleme kurulu, 2025-2026 programına sağladıkları işbirliği için aşağıda yer alan iş dünyasının değerli mensuplarına teşekkür eder.

The organizing committee thanks the esteemed company representatives listed below for their cooperation to run 2025-2026 program.

A101 Yeni Mağazacılık
Ayşegül Bekar
İhsan Eren
İsmail Gökhan Kaya

Beko Bulaşık Makinesi İşletmesi
Elshad Mahmudov
Beytullah Uyulur

Beko Buzdolabı İşletmesi
Hüseyin Çoker
Merve Özhavutçu
Özlem Deviren Uslu

Bilkent Holding
Ceyda Kitapçıgil

Emeklilik Gözetim Merkezi
Mustafa Koray Aytekin
Hicran Aktaş Coşar
Barış Çakmak
Ege Kaçar
Yasemin Yüksel Çuhadar Kütle
Samet Özben

Eti Gıda
Samet Aruca
Irmak Çelikkol Çalışkan
Aslı Demirgil Günhan
Fatih Yarımoğlu
Elif Nur Yılmaz

Hayat Finans
Özer Baran
Alperen Doğan
Suzan Öztürk

Karel
Arda Baş
Alper Sarıkan

Limak Çimento
Bensu Doğan Basmacı
Ayça Küçükoğlu

Mavi Jeans
Ensar Emirali

Nesco Gıda
Nazlı Esen Ummansu

Roketsan
Zeynep Alpaslan
Bilgehan Bıyıklı
Tevfik Buğra Ünsal
Selin Tamer

SCW.AI
Haluk Atlı

TeklifimGelsin
Hakan Biri
Alper Yenigün

Teknopar
Ertuğrul Bilgiç Kemal Ekbiç
İlknur Öz
Perin Ünal

Tepe Kurumsal Çözümler
Yakup Altay
Büşra Doğan Doysan
Melih Günaydın
Mert Can Taburlu
Uğur Ünlükalaycı

Türk Traktör
Murat Göçer
Elif Kaplan
Cem Tüfekçi

Türkiye İş Bankası
Oğuz G. Agun
Gülşah Bozdağ
Selçuk Coşkun
Tuğçe İlayda Nar

Unilever Türkiye
Nisanur Mugayıtoğlu
Bilgehan Tamkoç

Bölüm Başkanı'ndan

Bilkent Üniversitesi Endüstri Mühendisliği Bölümü, öğrencilerinin teknolojik ve sosyal değişikliklere uyum sağlayan, yaşam boyu öğrenen ve sorgulayan iyi endüstri mühendisleri olarak mezun olmalarını amaçlamaktadır. Karmaşık sistemlere ve problemlere bütün olarak bakabilme ve analitik düşünebilme, eğitim programının önemli amaçlarındandır. Bölüm, 2007 yılında *Accreditation Board for Engineering and Technology (ABET)* adlı bağımsız kuruluş tarafından eğitim kalitesini belgeleyen tam akreditasyonu Türkiye'de ilk alan mühendislik bölümüdür.

Eğitimde dünya çapında kalite standartlarını kullanan Endüstri Mühendisliği Bölümü, ülkemizde örnek gösterilen *Üniversite-Sanayi İşbirliği Programı*'nı 32 yıldır başarıyla uygulamaktadır. Programın hedefi mezuniyet aşamasındaki öğrencilere kapsamlı mesleki deneyim kazandırmaktır. Altı-yedi kişilik proje ekipleri, akademik ve endüstriyel danışmanların gözetiminde firmaların çözüm bekleyen gerçek problemlerini çözmektedirler.

Bu yıl, *24. Endüstri Mühendisliği Proje Fuarı ve Yarışması*'nda 21 proje bulunmaktadır. Fuarda öğrencilerimiz, yıl boyunca projeleri üzerinde yaptıkları çalışmalarını sunmaktadırlar. Onları özverili çalışmaları için kutluyor, programa büyük katkıları olan firma yetkililerine ve danışmanlarımıza teşekkür ediyorum.

Bütün süreç boyunca yoğun ve özverili çalışmalarıyla programın hedeflerine ulaşması için büyük çaba gösteren program koordinatörleri Prof. Dr. Savaş Dayanık, Doç. Dr. Emre Nadar, ve Dr. Emre Uzun'a, Üniversite-Sanayi İşbirliği Öğrenci Projeleri Koordinatörü'müz Yeşim Gülseren'e, asistanlarımız, Aslı Baytar, Semra Selin Eraslan, Fatih Selim Erdem, Ece Kuşdemir ve emeği geçen herkese çok teşekkür ediyorum.

Prof. Dr. Bahar Yetiş Kara
Endüstri Mühendisliği Bölüm Başkanı

Chairperson's Message

Bilkent University Industrial Engineering Department strives for its students to grasp changes in technology and society and be lifelong learners and inquirers. One of the department's educational goals is that our students hold a holistic view of systems and problems backed up with analytical thinking. The department is the first engineering department in Turkey, the quality of whose education program was fully accredited by *the Accreditation Board for Engineering and Technology (ABET)* back in 2007.

For 32 years, the Industrial Engineering Department has been successfully running its exemplary *University-Industry Collaboration Program.* The program's objective is to have the department's senior students gain full-fledged industrial experience before getting full industrial positions. Six-to-seven member student groups attack real open problems of companies under the supervision of academic and industrial advisors.

Twenty-one projects are present at the *24th Industrial Engineering Project Fair and Competition.* At the fair, student groups present their year-long work and the outcomes of their projects. I congratulate them for their tireless and heart-whole hard work. I also thank the company representatives and academic and industrial advisors for their support and collaboration.

Finally, I thank course coordinators Prof. Dr. Savaş Dayanık, Assoc. Prof. Emre Nadar, and Dr. Emre Uzun, University-Industry Collaboration Student Projects Coordinator Yeşim Gülseren, graduate assistants Aslı Baytar, Semra Selin Eraslan, Fatih Selim Erdem, Ece Kuşdemir for their relentless efforts to ensure that the program succeeds.

Prof. Dr. Bahar Yetiş Kara
Industrial Engineering Department Chairperson

Teşekkür Mektupları

Appreciation Letters

A101 olarak, Türkiye'nin dört bir yanındaki geniş mağaza ağımıza ürünlerimizi ulaştırırken operasyonel verimliliği ve dijitalleşmeyi odağımızda tutuyoruz. Operasyonel süreçlerimizdeki hızı artırmak ve kaynaklarımızı daha etkin kullanmak adına Bilkent Üniversitesi Endüstri Mühendisliği Bölümü ile kıymetli bir iş birliğine imza attık.

Bu kapsamda, bitirme projesi öğrencileri ve değerli danışman hocalarımızla birlikte depolarımızdaki ürün yerleşimini optimize eden bir çalışma gerçekleştirdik. Geliştirilen stratejik yerleşim planı sayesinde, operatörlerimizin depo içindeki yürüme mesafelerini minimize ederken, ortalama koli toplama sayılarımızda dikkate değer bir artış sağladık.

Genç mühendis adaylarımızın teorik bilgilerini gerçek saha verileriyle birleştirerek sundukları bu çözüm, operasyonel çevikliğimize büyük katkı sağladı. Bizlere sundukları yeni perspektifler ve titiz çalışmaları için Bilkent Üniversitesi Endüstri Mühendisliği Bölümü'ne, kıymetli hocalarımıza ve projede emeği geçen tüm öğrenci arkadaşlarımıza teşekkür ederiz.

İsmail Gökhan Kaya
Depo Operasyonları ve Projeler Müdürü

Emeklilik Gözetim Merkezi A.Ş. (EGM), 4632 sayılı Kanun uyarınca, Hazine ve Maliye Bakanlığı'nın görevlendirme ve yetkilendirmesi çerçevesinde 10 Temmuz 2003 tarihinde kurulmuştur. EGM'nin kanunla belirtilen amaçları, özel emeklilik sisteminin güvenli ve etkin biçimde işlemesini sağlamak, katılımcıların hak ve menfaatlerini korumaktır. Bu doğrultuda emeklilik şirketleri, fon yöneticileri ve aracılar üzerindeki gözetime konu altyapıyı sağlamakta, Sigortacılık ve Özel Emeklilik Düzenleme ve Denetleme Kurumu (SEDDK) ve Sermaye Piyasası Kurulu (SPK) başta olmak üzere ilgili otoritelerle yakın bir şekilde çalışmaktadır. Aynı zamanda özel emeklilik sisteminin performansını artırmaya yönelik öneriler üzerinde çalışmakta ve bu önerileri yetkili otoritelere sunmaktadır.

EGM çatısı altında faaliyet gösteren Dijital Ürün Yönetimi, Strateji, Sistem Geliştirme ve AR-GE ekibi; EGM'nin dijital dönüşümüne yön veren, kullanıcı odaklı ve veri temelli çözümler geliştiren birimidir. Bu birim, özel emeklilik sisteminde katılımcıların ihtiyaçlarını analiz ederek yenilikçi çözümler sunar. Aynı zamanda sistem performansını artırmaya yönelik otomasyon projeleri geliştirir, regülasyonlara uyumlu süreçler oluşturur ve küresel eğilimleri takip ederek sistem iyileştirme önerileri hazırlar. Ekip, kurumun dijital vizyonunu destekleyen stratejik kararların alınmasında da aktif rol oynar. EGM olarak üniversitelerde üretilen bilimsel bilginin, özel emeklilik sisteminin gelişimine katkı sağlayacak şekilde sektöre aktarılmasının hem katılımcı memnuniyetinin artırılmasına hem de sistemin sürdürülebilirliği açısından kritik bir öneme sahip olduğunun farkındayız. Bu önemin farkında olarak Emeklilik Gözetim Merkezi'nde üniversiteler ile iş birliği süreçlerini işletip bu süreçleri yakından takip ediyoruz. Bu sayede bilimsel yaklaşımın ve çok paydaşlı iş birliklerinin getirdiği değerin altını bu çalışmalarla güçlendirdiğimize inanıyoruz.

Bilkent Üniversitesi Endüstri Mühendisliği Sanayi İş Birliği Projeleri kapsamında bu yıl iki adet projeyi mühendis adayı öğrenciler ile tamamladık. "İllikit Menkul Kıymetleri Getiri Eğrisi ile Değerleme Modeli" projesi kapsamında, likit olmayan sabit getirili menkul kıymetlerin değerleme yöntemlerindeki tutarsızlıkların giderilmesi amacıyla Nelson-Siegel-Svensson (NSS) modeli geliştirilmiştir. Bu sayede, tüm menkul kıymetlerin ortak bir yöntemle değerlenmesi, maksimum getirinin elde edilmesi ve adaletsiz servet transferinin önlenmesi hedeflenmektedir. "Bireysel Emeklilik Sisteminde Devlet Katkısı Sistem Eniyilemesi" projesi kapsamında ise negatif reel getiri üreten ve performans farklılıkları nedeniyle katılımcılar arasında eşitsizliğe neden olan Devlet katkısı fonları incelenmiş ve geçmiş piyasa verileri çerçevesinde getiriyi maksimize eden, portföy yapılarının analizine imkân veren bir öneri modeli geliştirilmiştir. Devlet katkısı fonlarının portföy sınırlamasına ilişkin kısıtlamaların değiştirilmesine katkı sağlayacak bu çalışma, sistemin iyileştirilmesine yönelik önerilere önemli bir örnek teşkil etmektedir.

Bilkent Üniversitesi Endüstri Mühendisliği Sanayi İş Birliği Projeleri kapsamında mühendis adayı öğrencilerle birlikte gerçek sektör ihtiyaçlarını karşılamaya yönelik somut çıktılar üretmenin yanı sıra, öğrencilerin mesleki gelişimlerine doğrudan katkı sağlamak amacıyla yürüttüğümüz bu çalışmalar, Bilkent Üniversitesi Endüstri Mühendisliği Bölümü ile paylaştığımız ortak vizyonun güçlü bir yansımasıdır. Proje sürecinde öğrencilerin sergilediği kararlılık ve gayret, akademisyenlerin rehberliğiyle birleşerek yüksek katma değerli sonuçların elde edilmesini sağlamıştır. Bu süreçte ortaya çıkan sinerjiye katkı sunan başta Projeler Koordinatörlüğü olmak üzere tüm akademik kadroya ve projelerde görev alan tüm öğrencilere içten teşekkürlerimizi sunar; merak ve öğrenme isteklerinin daima canlı kalmasını, başarılarının ise yaşamları boyunca artarak devam etmesini dileriz.

Mustafa Akmaz
Genel Müdür

Hayat Finans

Hayat Finans, Türkiye'nin tamamı Türk sermayeli ilk dijital bankası olarak, güçlü teknolojik altyapısı ve kesintisiz hizmet anlayışıyla bireysel ve ticari bankacılıktan yatırıma kadar geniş bir yelpazede 2023 yılından bu yana yenilikçi finansal çözümler sunmaktadır. Güvenilir, hızlı ve müşteri odaklı yaklaşımıyla dijital bankacılık deneyimini yeniden tanımlayan Hayat Finans, sürekli gelişen teknolojiyi merkeze alarak kullanıcılarına hayatın içinde benzersiz bir finansal deneyim sağlamayı hedeflemektedir.

Hayat Finans olarak, şubesiz dijital bankacılık süreçlerinde otomatik karar alma mekanizmalarımızı destekleyen finansal modellerimizi veri odaklı yapmak bizim için kritik bir önem taşımaktadır. Bu bağlamda yaptığımız proje, Bilkent Üniversitesi tedrisatında yetişmiş, her biri bulunduğu yere katma değer sağlamak konusunda yüksek ümit vadeden arkadaşlarımızın analitik bakış açılarını gerçek hayat problemleri ile buluştururken aynı zamanda bizim için de farklı perspektifler kazanma imkanı sunmuştur.

Bankacılık dünyasının en temel problemlerinden olan müşterilerin gelirlerini tahmin etmek üzerine akademik yaklaşımla yeni bir model hazırladığımız bu proje, Bankamıza; müşteri segmentasyonu, doğru pazarlama faaliyetleri ve müşteriye isabetli tutarlarda kredi limiti hesaplamak gibi direkt ve somut katkılar sağlayacaktır.

Süreç boyunca bilgi ve deneyimleriyle projeye yön veren danışman hocamız Sayın Prof. Dr. Savaş Dayanık'a ve organizasyondaki katkıları için ÜSİ Mezuniyet Projeleri Koordinatörü Sayın Yeşim Gülseren Hanım'a içten teşekkürlerimizi sunarız.

Bu kıymetli iş birliğinde emeği geçen tüm öğrencilere ve Bilkent Üniversitesi'ne katkıları için teşekkür eder gelecekte de benzer projelerde yeniden bir araya gelmeyi temenni ederiz.

Özer Baran
Genel Müdür (V)

Haberleşme ve savunma teknolojileri başta olmak üzere elektronik üretim ve sistem entegrasyonu alanlarında faaliyet gösteren Karel Elektronik, yenilikçi üretim anlayışı ve güçlü Ar-Ge altyapısı ile sektöründe önemli bir konuma sahiptir. Gelişmiş üretim tesisleri, nitelikli insan kaynağı ve teknoloji odaklı yaklaşımı ile Karel, hem ulusal hem de uluslararası pazarlarda rekabetçi çözümler sunmaktadır.

Sanayi ile akademi arasındaki iş birliği, teknolojik gelişimin sürdürülebilirliği ve yenilikçi çözümlerin hayata geçirilmesi açısından büyük önem taşımaktadır. Bu doğrultuda yürütülen üniversite-sanayi iş birlikleri kapsamında gerçekleştirilen lisans bitirme projeleri, gerçek üretim problemlerine çözüm geliştirilmesine katkı sağlamakta ve mühendis adaylarının pratik deneyim kazanmalarına olanak tanımaktadır.

Bilkent Üniversitesi Endüstri Mühendisliği Bölümü öğrencileri tarafından yürütülen ve Karel Elektronik'in PCB üretim süreçlerine yönelik geliştirilen "Baskılı Devre Kartı Üretim Çizelgelemesi" projesi, üretim planlama ve çizelgeleme süreçlerinin daha verimli, veri odaklı ve sistematik bir şekilde yönetilmesine katkı sağlayacak nitelikte değerli çıktılar ortaya koymuştur.

Proje kapsamında geliştirilen karar destek sistemi ile üretim ortamının karmaşık yapısı dikkate alınarak makine atamaları, iş sıralamaları ve yeniden planlama süreçleri için uygulanabilir ve etkin çözümler sunulmuştur. Elde edilen sonuçların, üretim performansının artırılması ve teslimat güvenilirliğinin iyileştirilmesi açısından önemli katkılar sağlayacağı değerlendirilmektedir.

Bu değerli çalışmanın gerçekleştirilmesinde emeği geçen Bilkent Üniversitesi Rektörlüğü'ne, Mühendislik Fakültesi yönetici ve akademisyenlerine, özellikle proje sürecine akademik katkılarıyla yön veren danışman hocamıza, proje ekibinde yer alan kıymetli mühendis adaylarına ve süreç boyunca destek sağlayan tüm paydaşlara teşekkür ederiz.

Üniversite-sanayi iş birliğinin güçlenerek devam etmesini temenni ediyor, bu tür projelerin hem akademik hem de endüstriyel açıdan yüksek katma değer yaratmaya devam edeceğine inanıyoruz.

Arda Baş
Askeri Sistemler Proje Yöneticisi

Limak Çimento
now

Limak Holding iştiraki olarak, 2000 yılından bu yana çimento sektöründe faaliyet göstermekteyiz. 11 çimento fabrikamız, 30'dan fazla hazır beton tesisimiz ve 2000'den fazla çalışanımız ile sektördeki varlığımızı güçlendiriyor; güçlü dekarbonizasyon yol haritamız doğrultusunda, "Mevcut En İyi Teknikler"i benimseyerek, ileri teknolojiler ve sektördeki öncü adımlarımızla ilerlemeye devam ediyoruz. Yüksek hacimli üretim ve lojistik operasyonlarımız kapsamında süreç verimliliğini artırmak, operasyonel hataları minimize etmek ve sürdürülebilir çözümler geliştirmek öncelikli hedeflerimiz arasında yer almaktadır.

Bu doğrultuda, Bilkent Üniversitesi ile yürütülen Sanayi Odaklı Bitirme Projeleri kapsamında gerçekleştirilen "Kantar Optimizasyonu" çalışması, iş süreçlerimizin analiz edilmesi ve geliştirilmesine önemli katkılar sağlamıştır. Proje kapsamında gerçekleştirilen analizler ve geliştirilen simülasyon modeli sayesinde mevcut sistemdeki darboğazlar net bir şekilde ortaya konmuş; manuel veri girişinden kaynaklanan hatalar, operasyon süreleri ve süreç verimsizlikleri detaylı olarak incelenmiştir. Önerilen QR tabanlı veri aktarımı, otomatik belge üretimi ve fiziksel iyileştirmeler gibi çözümler, operasyonel performansımızın artırılması adına değerli çıktılar sunmuştur. Bu iş birliği sayesinde, hem şirketimiz gerçek operasyonel problemlerine yenilikçi ve uygulanabilir çözümler geliştirme fırsatı bulmuş, hem de mühendis adaylarımız gerçek saha deneyimi kazanarak iş hayatına hazırlık sürecinde önemli bir adım atmıştır. Akademik bilgi ile sanayi tecrübesinin bir araya gelmesiyle ortaya çıkan bu sinerjinin, gelecekte yapılacak çalışmalar için de yol gösterici olacağına inanıyoruz.

Bu değerli çalışma sürecinde katkı sağlayan Bilkent Üniversitesi Rektörlüğü'ne, Mühendislik Fakültesi akademisyenlerine, proje danışmanlarımıza ve emeği geçen tüm öğrenci arkadaşlarımıza teşekkür eder; üniversite-sanayi iş birliğinin artarak devam etmesini temenni ederiz.

Z. Ayça Küçükoğlu
Grup Lojistik Müdürü

1991 yılında İstanbul'da kurulan Mavi, 34 yıllık denim uzmanlığından aldığı güçle bugün global bir lifestyle markası konumunda. 2017'de halka açılan şirket, Türkiye, ABD, Kanada, Almanya ve Rusya'nın aralarında bulunduğu 34 ülkede, 498'i Mavi shop olmak üzere 4.000 noktada perakende, toptan ve online kanallar aracılığıyla müşterileriyle buluşuyor.

Mavi, All Blue stratejisi çerçevesinde, sürdürülebilirliği şirket kültürüne, vizyonuna, ürünlerine ve büyüme hedeflerine entegre etmeye devam ediyor. TIME ve Statista, 2026 yılının başında, Mavi'nin sürdürülebilir büyümede dünyanın en iyi 2., hazır giyimde ise 1. şirketi olduğunu açıkladı. Kalbi denim ile atan ve dünyanın en iyi, en inovatif jean'lerini geliştirmek için tutkuyla çalışan 5.957 kişilik global Mavi ekibi de insan, çevre, denim ve toplum odaklı bu değerleri odağına alarak en iyiye ulaşma yolculuğunda markayı geleceğe taşıyor. Mavi olarak, faaliyet gösterdiğimiz farklı alanlarda yarattığımız katma değeri artırmayı, verimlilik ve sürdürülebilirlik odağında çalışmalar yürütmeyi önceliklendiriyoruz. Bu doğrultuda, her fırsatta süreçlerimizi geliştirecek, yenilikçi ve dijital çözümleri değerlendirmeye önem veriyoruz. Bu bakış açısıyla hem endüstri hem de öğrenciler için önemli bir değer yaratan Bilkent Üniversitesi ile yollarımız kesişti.

Bilkent Üniversitesi Endüstri Mühendisliği Bölümü'nde yürütmekte olduğumuz bitirme projeleri kapsamında bu yıl; lokasyon ve ürün bazında kısa dönem talep tahmini geliştirilmesine yönelik veri odaklı bir proje gerçekleştirilmiştir. Proje süresince öğrenci ekipleri, gerçek şirket verisi üzerinde çalışarak veri analizi, veri temizleme ve düzenleme, özellik çıkarımı ve farklı tahminleme modellerinin araştırılması ve uygulanması süreçlerinde aktif rol almış; elde ettikleri çıktıları sistematik ve ölçeklenebilir bir yapıya dönüştürerek farklı mağaza ve ürün gruplarına uygulanabilir bir tahmin sistemi ortaya koymuşlardır.

Bu projeler aracılığıyla akademi ve sektör iş birliğinin somut bir örneği hayata geçirilmiş; öğrenciler gerçek iş problemleri ve veri setleri üzerinde çalışma, veri kaynaklı zorlukları deneyimleme ve analitik yetkinliklerini geliştirme fırsatı bulmuştur. Proje çıktıları sayesinde mevcut yaklaşımlara kıyasla tahminleme doğruluğunda anlamlı iyileşmeler gözlemlenmiş; aynı zamanda geliştirilen çözümlerin sürdürülebilir ve genişletilebilir yapısı ile şirket içi uygulamalara katkı sağlayacak değerli içgörüler elde edilmiştir.

Bu süreçte verdikleri değerli katkılar için başta Endüstri Mühendisliği Bölümü olmak üzere, çok kıymetli danışman hocalarımıza ve projeyi gerçekleştiren proje grubunda yer alan ve emeği geçen sevgili öğrenci arkadaşlarımıza içten teşekkürlerimizi sunarız.

Ensar Emirali
Veri Analitiği Kıdemli Müdürü

SCW.AI

SCW.AI olarak üretim verimliliği, operasyonel mükemmellik ve optimizasyon alanlarında teknoloji çözümleri geliştiriyoruz. Üretim tesislerinden toplanan verileri analiz ederek işletmelerin daha verimli, daha hızlı ve daha doğru kararlar almasını sağlayan sistemler sunuyoruz. Veri analitiği, yapay zeka destekli tahminleme, üretim planlama ve çizelgeleme optimizasyonu gibi alanlarda geliştirdiğimiz çözümlerle üretim süreçlerinin dijital dönüşümüne katkı sağlıyoruz. Amacımız, üretim sahasındaki operasyonları daha görünür, ölçülebilir ve optimize edilebilir hale getirerek şirketlerin rekabet gücünü artırmak.

Sanayi ve akademi iş birliğinin, geleceğin üretim sistemlerini geliştirmede kritik bir rol oynadığına inanıyoruz. Her sene Bilkent Üniversitesi Endüstri Mühendisliği Bölümü son sınıf öğrencileriyle gerçekleştirdiğimiz bitirme projeleri, bizim için yalnızca akademik çalışma değil; bilgi paylaşımının, ortak üretimin ve yenilikçi bakış açılarının değerli örnekleri olmuşlardır. Öğrencilerin teorik bilgilerini gerçek üretim problemlerine uygulaması ve çözüm üretme süreçlerinde gösterdikleri analitik yaklaşım, üniversite-sanayi iş birliklerinin ne kadar önemli olduğunu ortaya koymaktadır. Bu tür ortak çalışmalar, genç mühendis adaylarının sektöre hazırlanmasına katkı sağlarken, bizlere de yeni bakış açıları ve yenilikçi fikirler kazandırıyor. Geleceğin üretim teknolojilerini şekillendirecek genç yeteneklerle birlikte çalışmak bizim için son derece kıymetli.

Her yıl benzer projeler aracılığıyla üniversiteler ve öğrencilerle bir araya gelmekten büyük memnuniyet duyuyoruz. Geleceğin üretim dünyasını birlikte şekillendirmeye devam edeceğimize inanıyoruz. Emek veren tüm öğrencilerimize ve değerli akademisyenlerimize teşekkür ederiz.

Haluk Atlı
Ar-Ge Direktörü

teklifimgelsin

2020 yılında kurulan TeklifimGelsin, kullanıcıların finansal süreçlerini yönetebildiği, ihtiyaçlarına uygun teklifleri karşılaştırıp başvurularını tek bir yerden kolayca tamamlayabildikleri bir kişiselleştirilmiş finansal platformdur. Finansal süreçleri herkes için ulaşılabilir, şeffaf ve hızlı hale getirme vizyonuyla hareket eden şirketimiz, bugün iki milyondan fazla kullanıcıya hizmet vermektedir.

Bilkent Üniversitesi'nin öğrencilerini hayat boyu öğrenen, çözümleyici ve bağımsız düşünebilen bireyler olarak yetiştirme vizyonu, şirketimizin genç, dinamik yapısıyla ve teknolojik yenilikleri sürdürülebilir ve bilinçli şekilde takip etme amacıyla ortak bir zemin oluşturmaktadır. Bu doğrultuda, öğrencilere, üniversitemize ve şirketlere değerli kazanımlar sunan, Endüstri Mühendisliği Bölümü tarafından yürütülen bitirme projeleri kapsamında bu yıl "Bireysel Kredilerde Onay Olasılığı Tahminleme Modeli" konulu projemizi tamamladık. Proje sürecimiz, öğrencilerimizin azimli çalışmaları ve Bilkent Üniversitesi'nin akademik birikimi ve rehberliği sayesinde başarıyla ilerledi. Bu süreç ile öğrencilerimizin mezun olmadan iş dünyasını ve gerçek çalışma ortamlarını yakından deneyimlemesine katkı sağlamaktan büyük mutluluk duyuyoruz.

Projenin hayata geçmesinde emeği geçen ve bu kazanımların edinilmesine imkan sağlayan Endüstri Mühendisliği Bölümü'ne, saygıdeğer danışman hocalarımız başta olmak üzere süreç ile ilgilenen tüm akademik kadroya ve projelerde emeği geçen öğrenci gruplarına teşekkürlerimizi sunarız.

Öğrencilerimizin mezuniyetini şimdiden kutlar, gelecekteki başarılı çalışmalarında ve kariyer yollarında kendilerini takip etmekten mutluluk duyarız.

Alper Yenigün
Ürün Müdürü
Bilkent IE 2021 Mezunu

TÜRKİYE İŞ BANKASI

1924 yılında kurucumuzun önderliğinde ülkenin kalkındırılması ve ekonomimizin bağımsızlığını sürdürmesi için kurulan Bankamız, 100 yılın ardından ikinci yüzyılında da bu amaç doğrultusunda çalışmaktadır. Çalışanların üye olduğu İş Bankası Munzam Sandık Vakfı sayesinde Bankamızın en büyük ortağı, kendi çalışanlarıdır. Değişimin, gelişimin, adaptasyonun ve sürdürülebilirliğin çok önemli olduğu süreçlerimizde hem işin sahibi hem de en önemli kaynağı olan çalışma arkadaşlarımızla "Benim İşim" diyerek fayda yaratmak için çalışmayı sürdürüyoruz.

Bankamız, Türkiye'deki ilk ATM'yi (Bankamatik) 25 Aralık 1987 tarihinde Ankara'da hizmete sunmuştur. Yenişehir/Ankara Şubesi'nde kurulan bu cihazla Türkiye, 7/24 bankacılık hizmetiyle tanışmıştır. Bankamız, "Bankamatik" markasının tescilli sahibidir ve bu teknolojiyle Türk bankacılık sektöründe öncü olmuştur. Yaygın Bankamatik ağı ile özel Bankalar arasında ilk sırada yer almakta olup sağladığımız teknolojik geliştirmelerle kesintisiz hizmet sağlama hedeflerimiz doğrultusunda ilklerin Bankası olarak çalışmalarımızı sürdürüyoruz.

Kaynağı üniversitelerde olan bilimsel bilginin sanayinin Teknoloji geliştirme çalışmalarına aktarılmasının çok kıymetli olduğu inanıyoruz. Hedeflerimiz doğrultusunda yeni teknolojileri üretmek, üretilenlere paydaş olmanın yanı sıra optimizasyon ve dijital verimlilik fırsatlarını her daim değerlendiriyoruz. Bu noktada ise yollarımız hem endüstri hem de öğrenciler için büyük katma değer oluşturan Bilkent Üniversitesi ile kesişti. Bilkent Üniversitesi Endüstri Mühendisliği bölümünde yürütmekte olduğumuz bitirme projeleri kapsamında Bankamatik cihazlarımızın verilerini işleyerek olası arızalarının öncesinde tahminleme modelini değerli öğrenci gruplarımız ve danışman hocalarımızın destekleri ile tamamladık

Proje kapsamında bir araya geldiğimiz, emekleri ile süreçlerimize destek olmuş hedeflenen çıktı ve kazanımlara ulaşmamızda katkı sağlamış öğrenci arkadaşlarımıza, görüşleriyle kıymetli arkadaşlarımıza yol gösteren Bilkent Üniversitesi Akademisyenleri'ne ve süreçleriyle bizi bir araya getiren Üniversite- Sanayi İşbirliği Koordinatörleri'ne çok teşekkür ederiz.

Selçuk Coşkun
Müşteri İlişkileri Bölümü
II. Müdür

TürkTraktör

Türkiye'nin üretim öncülerinden TürkTraktör, endüstriyel hayatına 1954 yılında başlamış ve bugün gelinen noktada 125 ülkeye yaptığı ihracatla ülkemizin önemli kuruluşlarından biri olmuştur. Türk Traktör'ün kurumsal yapılanmasını, Koç Holding ve CNH Industrial ortaklığı oluşturmaktadır.

Türk Traktör, New Holland ve Case IH marka traktörleri aynı anda üretebilme yetkinliğine sahiptir, 2020 yılından beri yerli kazıcı yükleyici üretimine de devam etmektedir.

Üretiminde göstermiş olduğu hassasiyeti satış ve satış sonrasında da titizlikle devam ettiren Türk Traktör, Türkiye genelindeki bayi ağı ve servisleriyle çiftçinin güvenilir yoldaşı olma misyonunu pekiştirmektedir.

Türk Traktör olarak, tüm süreçlerimizde verimlilik, sürdürülebilirlik ve sürekli iyileştirme bakış açısını önceliklendiriyor; optimizasyon ve dijitalleşme fırsatlarını yakından takip ediyoruz. Bu yaklaşımımızın önemli bir parçası olarak, akademi-sanayi iş birliğine büyük değer veriyoruz. Bu kapsamda Bilkent Üniversitesi Endüstri Mühendisliği Bölümü ile "Sevkiyat programları ile uyumlu dağıtım planlama karar destek sistemi" projesi üzerinde çalıştık.Gerçek iş problemleri üzerinden ilerleyen bu projeler sayesinde öğrencilerimiz, teorik bilgilerini pratikle buluşturma ve iş hayatına daha donanımlı şekilde hazırlanma fırsatı yakalarken; bizler de süreçlerimize dış bir gözle bakma, gelişim alanlarımızı daha net görme ve yeni bakış açıları kazanma imkânı elde ettik. Aynı zamanda akademik dünyadaki güncel yaklaşımları takip ederek, sürekli gelişim kültürümüzü daha da güçlendirdik.

Bu değerli iş birliği için Bilkent Üniversitesi Endüstri Mühendisliği

Bölümü'ne, kıymetli danışman hocalarımıza ve projelerde emeği geçen tüm öğrenci arkadaşlarımıza teşekkür eder; ilerleyen dönemlerde de benzer iş birliklerini sürdürmeyi dileriz.

Murat Göçer
Satış Sonrası Tedarik Zinciri Depo ve Dağıtım Yön.
Ürün Lideri

Dünyanın en büyük hızlı tüketim ürünleri şirketlerinden birisi olan Unilever olarak, Türkiye'de 100 yılı aşkın bir süredir faaliyet gösteriyoruz. Güçlü ve amaç sahibi markalarımız ve sorumlu iş yapmanın üstün ve sürdürülebilir performans sağladığına olan inancımızla operasyonlarımızı sürdürüyoruz. Ürünlerimiz 190 ülkede 4.4 milyon perakende noktasında her gün 3.4 milyar insana ulaşıyor. Dünya üzerinde 280'den fazla, Türkiye'de ise 2 üretim fabrikası ile Türkiye'nin en gözde şirketleri arasında yer alıyoruz. "Sürdürülebilir yaşamı mümkün kılmak" ve "Hep birlikte daha iyi bir dünya için" temalarını tüm çalışmalarımıza temel alıyoruz.

Unilever ailesi olarak tedarik zincirinde iş güvenliğine, her zaman "Sıfır Tolerans" ilkesi ile yaklaşıp, birlikte çalıştığımız tüm arkadaşlarımızın güvenliğini, esenliğini her geçen gün ileriye taşımak için sürekli olarak çalışıyoruz. Bu proje kapsamında fabrikamızda toz üretim bölümünde tecrübeye dayalı yapılan üretim planlamalarının, bilimsel bir metoda dayalı, kullanıcı dostu bir arayüz ile yapılmasını amaçladık. Çok kıymetli akademisyenlerimizin, akademik koordinatörümüzün destekleri ve yönlendirmeleri ile öğrencilerimiz firmamız için sezgisel metotlara dayanan bir algoritma geliştirdiler ve son derece karmaşık bir problemi modellemeyi başardılar. Bu modeli kullanıcı dostu bir arayüz ile tüm çalışanların kullanabileceği bir basitlikte kullanımımıza sundular. Proje vasıtası ile üretim tamamlama sürelerimizi iyileştirme ve işçilik maliyetlerimizi azaltma fırsatı bulduk. Projenin rakamsal çıktıları yanında projenin her aşamasında gerek zaman yönetimi gerek kuvvetli iletişimi ile Bilkent disiplinini hissettik. Bu disiplinin, bu karmaşık projenin sıkışık bir takvimde sonuca ulaşmasında önemli bir etken olduğunu gözlemledik.

Birlikte gerçekleştirdiğimiz bu katma değerli projenin Bilkent Üniversitesi'nde bugün okuyan ve gelecekte okuyacak tüm mühendis adaylarına yapacakları çalışmalarda önemli bir literatür katkısı olacağını umuyor, projede emeği geçen tüm akademisyenlerimize, koordinatörlerimize ve öğrencilerimize teşekkür ediyoruz. Tüm öğrencilerimize bundan sonraki iş ve akademik yaşamlarında başarılar diler, bir gün aramızda görmekten mutluluk duyarız.

Bilgehan Tamkoç
Toz Üretim ve BIH TR Projesi Kıdemli Müdürü

All the best of success...

and all the best of luck!

PROJELER
PROJECTS

Bankamatikler için Öngörücü Bakım Sistemi | 1

Türkiye İş Bankası

Proje Ekibi
Muhammed Hüseyin Beyribey, Alkım Bümen, Süleyman Burak Canıgür, Fatih Karaca, Halil Yaman Sağlık, Pelin Yalıncak, Dilara Büşra Yörür

Şirket Danışmanları
Oğuz G. Agun ve Selçuk Coşkun
İnsan Kaynakları Departmanı ve Bankamatik Birimi

Akademik Danışman
Prof. Dr. Savaş Dayanık
Endüstri Mühendisliği Bölümü

Özet

Bu projede kritik donanım arızalarını önceden tahminleyerek bankamatik çalışma sürelerini artırmak ve bakım kaynaklarını daha etkili biçimde yönlendirmek amaçlanmıştır. Bu doğrultuda arıza kayıtları ve işlem verileri birleştirilip günlük bazlı bir veri yapısı oluşturulmuş ve yakın gelecekteki arızaları tahminlemek için bir sınıflandırma modeli inşa edilmiştir. Tahmin çıktıları daha sonra bir eniyileme modeli ile birleştirilerek bir Karar Destek Sistemi oluşturulmuştur. Elde edilen sonuçlar doğrultusunda Türkiye geneli bankamatik çalışma sürelerinin yıllık %0,82 artışıyla birlikte milyarlarca liralık potansiyel işlem hacmi oluşacağı saptanmıştır. Geliştirilen tahminleyici bakım sistemi, Türkiye İş Bankası'nı yapay zeka tabanlı bankacılık operasyonlarında Türkiye'de ve küresel ölçekte sektör öncülerinden biri konumuna taşımaktadır.

Anahtar Sözcükler: Tahminleyici bakım, Bankamatik arıza tahmini, Karar destek sistemi, Makine öğrenmesi, Eniyileme

A Predictive Maintenance System Design for ATMs

Abstract

This project aims to increase ATM uptimes by forecasting critical hardware failures and directing maintenance resources more effectively. To this end, error logs and transaction data were integrated into a daily time-series structure, and a classification model was developed to predict failures in the near future. The prediction outputs were then integrated with an optimization model that selects priority ATMs under capacity constraints to create a Decision Support System. The results demonstrate a 0.82% annual increase in ATM uptime across Türkiye, unlocking a potential transaction volume of billions of Turkish Liras. The developed predictive maintenance system positions Türkiye İş Bankası as a global pioneer and the domestic leader in AI-driven banking operations.

Keywords: Predictive maintenance, ATM failure prediction, Decision support system, Machine learning, Optimization

1.1 Introduction

ATMs are one of the most visible service channels in banking, and even short periods of unavailability can directly affect customer experience. In a large ATM network, maintenance decisions therefore matter not only from a technical perspective but also from an operational one. Delays in intervention may reduce service continuity, while poorly prioritized maintenance may lead to inefficient use of already limited field resources (Agun, 2025; Türkiye İş Bankası, 2025).

1.1.1 Problem Definition

In current practice, ATM maintenance operations largely consist of reactive and preventive interventions. While such an approach can manage failures after they occur, it is less effective at ensuring service continuity before a disruption happens. This stems from the fact that maintenance capacity is limited, while failure risk is distributed across a large ATM network with varying usage patterns. Consequently, the challenge is not merely identifying potential failures but determining which ATMs should be prioritized (Agun, 2025; Türkiye İş Bankası, 2025).

1.1.2 Project Objective and Scope

The main objective of the project is to support predictive maintenance by detecting critical hardware failures before they occur and helping planners allocate maintenance resources more effectively under operational constraints. To achieve this, the project combines a machine learning based prediction layer with an optimization based decision support layer. The prediction component estimates whether an ATM is likely to experience a failure within a future horizon, while the decision support component prioritizes candidate ATMs under maintenance capacity constraints.

The scope of the study includes data processing, prediction model development, and an optimization-based decision support system (DSS) design that brings these components together in an interpretable and usable form as the main deliverable of the project.

1.2 Company and Process Background

The proposed solution is intended to support the existing operational structure of Türkiye İş Bankası's ATM division rather than replace it.

1.2.1 Company Background

Türkiye İş Bankası is one of the leading banking institutions in Turkey and TRNC, and operates a large nationwide ATM network as part of its retail banking infrastructure. For the bank, ATMs are not only physical service points but also an important interface between customers and everyday banking operations. From an operational perspective, the scale of the ATM network makes maintenance planning a nontrivial task. Machines are distributed across many locations with different usage intensities, transaction profiles, and local operational conditions (Türkiye Bankalar Birliği, 2024; Agun, 2025).

1.2.2 Current Process

Before the development of the proposed system, ATM maintenance decisions were primarily shaped by preventive and reactive interventions, where failures are typically handled after they occur. The current system may lead to downtimes exceeding 24 hours.

Not every potentially risky ATM can be inspected proactively. Error records capture technical events over time, whereas transaction data represent business activity at a different temporal resolution. Consequently, although valuable information exists, these raw inputs are difficult to use directly for maintenance prioritization (Agun, 2025).

1.2.3 Needs Analysis

The analysis of the current system indicated that the primary need was a decision-oriented structure capable of predicting failures and supporting maintenance prioritization. In other words, the company required a system that could transition from observing disruptions after they occurred to identifying which ATMs were most likely to require intervention in the near future. A second need concerned operational usability. Even a strong prediction model would have limited practical value if its outputs could not be translated into concrete maintenance actions. For this reason, the company context called for a system that could not only estimate short-term failure risk but also help determine which ATMs should be prioritized under categorical importance and capacity constraints.

1.3 Methodology

This section includes data, method, model, and system design procedures.

1.3.1 Data Collection and Preparation

The primary data sources consist of error logs recorded by ATM hardware and transaction records aggregated at the location level.

Before any modeling, the raw error logs undergo a cleaning pipeline. Duplicate records arising from overlapping data exports are removed, consecutive identical errors occurring within 60 seconds are merged into a single event, and short-lived system-level events are filtered out using predefined duration thresholds.

After cleaning, the system categorizes the remaining logs into four channels to filter noise and highlight actionable signals. The Target channel captures critical hardware faults. The Feature channel includes ten leading indicator signals that tend to precede critical faults. The Context channel tracks resource-state events. The Supervisor channel records historical maintenance interventions. The correlations across these channels motivate the use of a machine learning model to predict Target errors (Rosati et al., 2025).

To increase prediction performance, two additional channels are derived from the transaction records. Withdrawal and deposit counts are distributed to daily estimates using calendar-aware weights that reflect real usage patterns and then log-transformed to stabilize variance, forming the final six-channel input structure together with the four error-based channels. These six channels are: Target, Feature, Context, Supervisor, daily withdrawals, and daily deposits.

To align different data granularities, a Spatio-Temporal Spine framework is constructed by generating a Cartesian product of all active ATM serial numbers and all calendar dates in the observation period. The six channel values are left-joined onto this dense daily grid, with days without recorded events explicitly filled with zeros rather than left as missing values, allowing the model to detect signal density changes and degradation patterns more effectively than simple event counters.

1.3.2 Analysis and Model Development

Using the spine as input, the problem is formulated as a binary classification task. A sliding window of 31 days is advanced across each ATM's timeline to predict the probability of a critical failure within a seven-day horizon, ensuring the system focuses on recent operational behavior rather than stale data.

The feature extraction strategy draws on the observation that convolutional kernels in time-series classifiers such as HYDRA implicitly learn to detect statistical shape properties—trends, variance shifts, and temporal concentrations—within fixed-length windows (Dempster et al., 2023). Rather than learning these kernels end-to-end, the system extracts an equivalent set of shape statistics deterministically, preserving full interpretability and avoiding any model-fitting step for the representation layer while retaining the discriminative power that motivates convolutional approaches. This is particularly valuable in a banking context, where regulatory and audit requirements demand that every input to a maintenance decision be traceable and explainable.

Accordingly, the system extracts a deterministic set of 57 statistical shape features per window, organized into nine functional groups across six channels. Rather than treating each channel as a simple event counter, the feature design captures the shape of each signal over the 31-day window—its trend direction, variability, recency, and concentration—because two ATMs with identical error counts can exhibit very different risk profiles depending on whether those errors are accelerating, clustered near the present, or spread uniformly. The nine groups include mean activity level, standard deviation, global linear trend, early- and late-window slopes, Exponentially Weighted Moving Average, peak-to-mean ratio, Gini Coefficient, and Temporal Centroid. Three additional cross-channel features are used: Inverse Recency, a Supervisor Reset Effect indicator, and a Lead-Lag Ratio measuring how much Feature-channel activity precedes Target-channel failures.

Among the classifiers evaluated, XGBoost with a 31-day look-back window produced the best results. Hyperparameter search is conducted using Randomized-Search-CV with Time Series Split cross-validation to preserve chronological ordering and prevent data leakage. Class imbalance is handled through the scale-pos-weight parameter. To ensure that predicted probabilities reflect real failure frequencies, Isotonic Regression calibration is applied to a temporally held-out calibration set normalized using the training set's statistics (Zhang et al., 2016).

The effectiveness of the feature design is validated through SHAP feature importance analysis of the trained XGBoost model, which confirms that Inverse Recency and trend-based features are the most significant indicators of ATM failures. In particular, Inverse Recency captures short-term failure dependency, where recently failed machines are more likely to fail again, while trend-based features such as global and window-specific slopes enable the model to detect accelerating degradation over time. This alignment between the HYDRA-motivated feature design and the empirical importance ranking supports the choice of deterministic shape extraction over raw event frequencies.

The resulting risk scores are then used in a Deterministic Equivalent of a Stochastic Partitioned Knapsack Model that maximizes the Expected Protected Demand (Z) while considering regional technician capacities (Kellerer et al., 2004). The model is formulated as follows:

Sets and Indices

- J: Set of all locations, indexed by j.
- I_j: Set of all ATMs located within j, indexed by i.

Parameters

- p_{ij}: The failure probability for ATM i in province j.
- d_{ij}: The forecasted transaction volume for ATM i in province j.
- w_i: The category priority weight assigned to ATM i based on its operational classification.
- C_j: The maintenance capacity limit for province j.
- ρ: The minimum amount of money worth recovering.
- τ: The minimum probability threshold for eligibility, determined as 0.3425, the optimal threshold for Youden's J statistic.
- a_{ij}: Binary eligibility parameter used to linearize the threshold logic.

$$a_{ij} = 1 \Longleftrightarrow (p_{ij} \cdot d_{ij} \geq \rho) \wedge (p_{ij} \geq \tau)$$

Decision Variable

- x_{ij}: Binary variable equal to 1 if a maintenance team is dispatched to ATM i in province j, and 0 otherwise.

Objective Function

$$\max Z = \sum_{j \in J} \sum_{i \in I_j} (p_{ij} \cdot d_{ij} \cdot w_i) \cdot x_{ij} \tag{1.1}$$

Constraints

$$\sum_{i \in I_j} x_{ij} \leq C_j, \quad \forall j \in J \tag{1.2}$$

$$x_{ij} \leq a_{ij}, \quad \forall j \in J, \forall i \in I_j \tag{1.3}$$

$$x_{ij} \in \{0, 1\}, \quad \forall j \in J, \forall i \in I_j \tag{1.4}$$

Constraint (1.2) represents the regional capacity constraint for every location, whereas (1.3) enforces the threshold logic based on minimum probability and transaction volume. Constraint (1.4) is a simple binary dispatching constraint.

1.3.3 System Design

The proposed solution is delivered as a centralized, Python-based DSS that integrates predictive analytics with a graphical user interface (GUI; Figure 1.1). The system follows a “Predict-then-Optimize” workflow: the back-end engine generates failure probabilities for a seven-day horizon and then processes these through the partitioned Knapsack model to produce a feasible maintenance schedule.

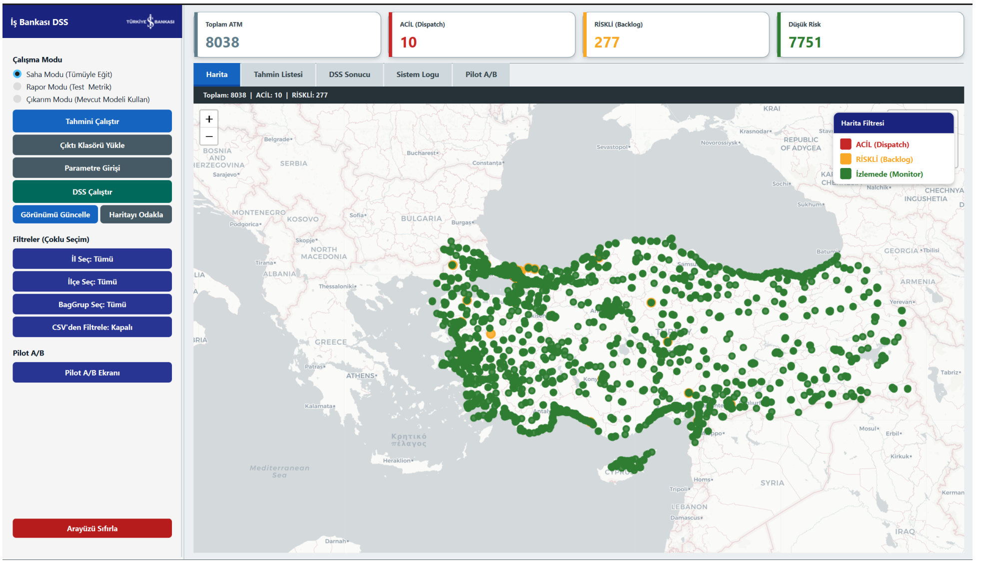

Figure 1.1: DSS Interface

The DSS includes an interactive dashboard for maintenance planners. A dynamic sidebar allows users to upload at least 31 days of historical transaction and error logs to refresh the model. The system is designed to run weekly in line with the bank's reporting cycles and maintenance logistics.

The dashboard summarizes network health by categorizing ATMs into three statuses: Dispatch (immediate visit), Backlog (at-risk machines exceeding regional capacity), and Monitor (stable machines). These results are visualized through a provincial risk map and detailed list views in the Prediction and DSS tabs for maintenance decisions. Through the sidebar, users can select the data to upload, working mode, and location- and CSV-based ATM filters, as well as adjust key parameters.

1.4 Implementation

The implementation is designed to integrate with Türkiye İş Bankası's existing maintenance ecosystem. The system operates as an iterative pipeline that uses updated error logs and transaction data to transform evolving risk profiles into maintenance schedules. After Türkiye İş Bankası installs the developed DSS application in its Cash Management Centers, the bank will continue its operations on a weekly basis, relying on the generated failure predictions and optimized maintenance priority results.

To integrate and evaluate the proposed system, a two-week pilot study is scheduled between the 6th of April and the 19th of April, which involves A/B testing. A-group ATMs receive a maintenance team on alert, whereas B-group follows the current maintenance schedule. This pilot study tests the developed system in a real environment and demonstrates how the predictive maintenance system can support the current maintenance system before broader deployment.

1.5 Results and Evaluation

This section includes the theoretical results and performance evaluation of the predictive maintenance system model.

1.5.1 Performance Results

The predictive maintenance system was theoretically tested for November 2024. Based on historical error logs, demand distributions, and mean time to repair (MTTR), an uptime increase of 0.82% was calculated, which exceeded the industrial advisor's expectation of 0.5% (Agun et al., 2026). This increase corresponds to hundreds of thousands of maintenance hours and billions of liras in protected transaction volume. In addition, the system is designed to improve customer satisfaction and operational efficiency. Customers would need to locate alternative ATMs less frequently, while the bank would be able to allocate maintenance resources more effectively for failures that were not predicted.

1.5.2 Validation

The model was validated using a date-based rolling window: training on historical data and testing on subsequent held-out periods. A 1,000-iteration bootstrap validation yielded a stable AUC of 0.7913 ± 0.0008, confirming robust discriminative power regardless of threshold choice or data split. At the threshold selected for deployment, roughly four out of five ATMs that actually failed were flagged in advance, and more than half of the flagged ATMs did experience a failure. Since missing a genuine failure risks over 24 hours of downtime while a false alarm wastes only one visit, the bank chose to favour broader coverage over dispatch efficiency—a balance adjustable at any time through the DSS interface. Industrial advisors confirmed that the resulting dispatch volume remains within operational capacity.

A pre-pilot study in March 2025 showed improved dispatch accuracy under real conditions, validating readiness for implementation. Experts from Türkiye İş Bankası further verified through face validity that the demand proxies reflect real-world customer behavior and maintenance schedules (Agun et al., 2026).

1.5.3 Discussion

The project theoretically yields a slightly higher performance increase than Türkiye İş Bankası's expectations. Both industrial and academic advisors confirmed that the project and implementation phase are feasible. Adopting such a system would represent one of the first predictive maintenance implementations for ATM networks in the banking sector (Agun, 2025; Agun et al., 2026).

1.6 Conclusion and Recommendations

This project developed an end-to-end predictive maintenance DSS for Türkiye İş Bankası. By combining HYDRA-motivated shape feature extraction with XGBoost classification and isotonic probability calibration, the system identifies the majority of ATMs likely to fail within a seven-day horizon while keeping unnecessary dispatches within operational capacity a balance the bank can adjust at any time through a configurable threshold. The estimated uptime increase is 0.82%, which represents a potential transaction volume increase of billions of Turkish Liras. The developed predictive maintenance system positions Türkiye İş Bankası as a pioneer in AI-driven banking operations.

With more detailed data, predictions could target specific hardware components rather than general failure risk. While the current model relies on transaction patterns and historical error logs, integrating real-time sensor data could further improve ATM uptime.

Bibliography

Agun, O. G. (2025). Kick-off Meeting with Türkiye İş Bankası (Zoom). Personal communication.

Agun, O. G., E. Bayramoğlu, and S. Ateş (2026). Progress and Pilot Study Meeting with Türkiye İş Bankası (Zoom). Personal communication.

Dempster, A., D. F. Schmidt, and G. I. Webb (2023). Hydra: Competing Convolutional Kernels for Fast and Accurate Time Series Classification. *Data Mining and Knowledge Discovery 37*(3), 1779–1805.

Kellerer, H., U. Pferschy, and D. Pisinger (2004). *Knapsack Problems.* Berlin, Heidelberg: Springer. Chapters 10 and 14.

Rosati, R., L. Romeo, and A. Mancini (2025). Single- and multi-task linear models for atms fault classification in human-centered predictive maintenance. *Computers & Industrial Engineering 200*, 110763.

Türkiye Bankalar Birliği (2024, may). Banka, Çalışan ve Şube sayıları. https://www.tbb.org.tr/sites/default/files/raporlar/51338/ekler/Banka_Calisan_ve_Sube_Sayilari-Mart_2024_0.pdf.

Türkiye İş Bankası (2025). İş Bankası Türkiye'nin En Büyük Özel Bankası Konumunu Sürdürüyor. https://www.isbank.com.tr/bankamizi-taniyin/is-bankasi-turkiyenin-en-buyuk-ozel-bankasi-konumunu-surduruyor.

Zhang, K., J. Xu, M. R. Min, G. Jiang, K. Pelechrinis, and H. Zhang (2016). Automated IT System Failure Prediction: A Deep Learning Approach. In *2016 IEEE International Conference on Big Data (Big Data)*, pp. 1291–1300. IEEE.

Baskılı Devre Kartı Üretim Çizelgelemesi | 2

Karel

Proje Ekibi
Mert Bulut, Alper Tuna Güçlü, Selin Haberal,
Erdem Hayfavi, Egemen İlhan, Ayşe Beyza Korkmaz, Melih Ulu

Şirket Danışmanı
Arda Baş
Askeri Sistemler Proje Yöneticisi

Akademik Danışman
Dr. Emre Uzun
Endüstri Mühendisliği Bölümü

Özet

Bu proje, Karel Elektronik'in baskılı devre kartı üretim hattı için veri temelli bir üretim planlama ve çizelgeleme sistemi geliştirmeyi amaçlamaktadır. Mevcut sistemde planlama büyük ölçüde Excel ve doğrudan iletişim yoluyla yürütülmekte, bu durum özellikle dizgi aşamasındaki yüksek kurulum süreleri nedeniyle kapasitenin etkin kullanılamamasına ve teslimat gecikmelerine yol açmaktadır. Geliştirilen yaklaşım; lehim pastası baskı, dizgi ve fırın aşamalarını bütünleşik biçimde ele almakta; işlem süreleri, kurulum süreleri, makine farklılıkları, tepsi bazlı parti yapısı, termal profil uyumluluğu ve teslim tarihlerini birlikte değerlendirmektedir. Çözüm kapsamında küçük örnekler için bir karma tamsayılı programlama modeli, gerçek boyutlu veriler için ise ATCS tabanlı bir sezgisel yöntem geliştirilmiştir. Tarihsel verilerle yapılan karşılaştırmalarda toplam ağırlıklı gecikmede yaklaşık %8.6 iyileşme ve zamanında teslim oranında %9.6 artış elde edilmiştir.

Anahtar Sözcükler: Üretim çizelgeleme, hibrit akış tipi atölye, toplam ağırlıklı gecikme, tabu arama, karar destek sistemi

Printed Circuit Board Production Scheduling

Abstract

This project develops a data-driven production planning and scheduling system for Karel Elektronik's printed circuit board manufacturing line. In the current system, planning is largely carried out manually through spreadsheets and direct coordination, which leads to frequent replanning, inefficient capacity usage, and delivery delays, especially at the pick and placement stage where setup times are substantial. The proposed approach models solder paste printing, pick and placement, and reflow oven operations in an integrated framework that considers processing times, setup times, machine differences, tray-based batching, thermal profile compatibility, and due dates simultaneously. The solution strategy combines an exact mixed-integer linear programming model for small instances with a heuristic for real size instances. The heuristic uses an ATCS-based initial sequence, a decoder that simulates the three-stage system, oven run construction based on profile compatibility, and a tabu search improvement phase. Historical company data were used to evaluate the approach. The results indicate an improvement of about 8.6% in total weighted tardiness and an increase of about 9.6% in on-time delivery performance compared with the current planning practice. Overall, the study provides a practical and explainable decision support tool for production planners.

Keywords: Production scheduling, hybrid flow shop, weighted tardiness, tabu search, decision support system

2.1 Company and Engineering Problem

Karel Elektronik is a major Turkish manufacturer operating in communication, defense, automotive, field operation technologies, and electronic manufacturing services. Within its electronic manufacturing services activities, the company performs printed circuit board (PCB) assembly, prototyping, automated testing, and system integration. The production environment considered in this project is a three-stage PCB assembly line consisting of solder paste printing, pick and placement, and reflow soldering. This structure naturally leads to a multi-stage scheduling problem in which jobs compete for parallel resources and are affected by both setup requirements and due dates. Similar production environments are commonly modeled as hybrid flow shops in the scheduling literature (Pinedo, 2016; Oujana et al., 2019).

In the current practice, once a work order is released, material availability is checked in the ERP system and production is scheduled largely by planners using spreadsheets and direct communication. Although the ERP system supports inventory and order management, it does not provide an integrated scheduling logic that simultaneously accounts for machine capacities, setup dependencies, tray-based batch sizes, and oven compatibility rules. As a result, the plan is revised frequently during the day due to material shortages, urgent jobs, and machine-related issues.

The most critical engineering problem is the lack of a systematic scheduling mechanism that can reduce late deliveries while preserving feasibility on the shop floor. In particular, the pick and placement stage has long sequence-dependent setup times, and some sequencing decisions lead to significant capacity losses. Therefore, the project focuses on designing a decision support system that generates feasible schedules while minimizing delivery-related performance losses.

The main performance measures are total weighted tardiness and on-time delivery rate. The main deliverables of the project are a production scheduling model, a Python-based solution engine, spreadsheet-based input and output interfaces, and a rescheduling logic suitable for daily operational use.

2.2 Current System Analysis

The production line examined in this study follows a fixed technological sequence. First, PCBs are loaded onto trays and processed at the solder paste printing stage. Second, they are sent to pick and placement machines where electronic components are mounted. Finally, they enter the reflow stage, where soldering is completed in ovens. Once boards are loaded into trays, batch composition remains fixed throughout the line.

The company operates multiple machines at the first two stages and a limited number of reflow ovens at the last stage. Although all products can be processed on all machines, machines differ in speed and technical characteristics. In addition, setups depend on product changes. Setup times at solder paste printing and reflow are relatively short, while pick and placement setups can vary from a few hours to much longer durations depending on feeder, stencil, nozzle, and program changes. This makes sequencing particularly important.

A key observation from the on-site study is that products requiring similar setups are not always processed consecutively. Since manual planning cannot fully exploit product family similarities, setup frequency increases and effective capacity decreases. Moreover, the reflow stage adds another layer of complexity because batches with compatible temperature profiles may be processed together in the same oven run. The current planning process does not evaluate this trade-off systematically.

These observations indicate that the scheduling problem cannot be handled adequately by simple first-come-first-served or purely manual rules. It requires a model that jointly considers batching, parallel machine assignment, sequence-dependent setups, precedence relations among stages, and oven grouping decisions.

2.3 Model and Proposed System

The proposed system is designed as a decision support tool that converts work-order data into feasible and interpretable production schedules. The solution framework has three layers: data preparation, scheduling, and output generation.

2.3.1 Modeling Framework

The production environment is modeled as a hybrid flexible flow shop with three sequential stages and parallel machines at each stage. Jobs correspond to batches rather than individual boards. For each batch, the system determines machine assignments, stage-wise start and completion times, and sequencing decisions. At the reflow stage, the model also determines which compatible batches should be grouped in the same oven run. This abstraction is consistent with standard scheduling formulations discussed in the literature (Pinedo, 2016; Oujana et al., 2019). The main assumptions are

- jobs are not preempted once processing starts,
- processing times, setup durations, due dates, and release dates are assumed known,
- transfer times between stages are neglected,
- intermediate buffer capacity is assumed sufficient,
- rework is ignored in the baseline model,
- oven grouping is allowed only for batches with compatible thermal profiles.

The primary objective is to minimize total weighted tardiness. Let R_w denote the completion time of work order w, d_w its due date, π_w its penalty weight, and Z_w its tardiness.

The notation used in the mathematical model is summarized here for completeness.

Sets and indices.

- $\mathcal{W}$: set of work orders, index w.
- $\mathcal{T}$: set of PCB types, index t.
- $\mathcal{P} = \{1, 2, 3\}$: set of processes, index p.
- $\mathcal{M}_p$: set of machines at process $p \in \mathcal{P}$, index m.
- $\mathcal{B}_{wt}$: set of potential batches for type t in work order w, index b.
- $\mathcal{B} = \{(w, t, b) \mid w \in \mathcal{W},\ t \in \mathcal{T},\ b \in \mathcal{B}_{wt}\}$: set of all batches.

Parameters. All time-related parameters are measured in minutes.

- n_{wt}: number of cards demanded of type t in work order w.
- cap_t: batch capacity for type t.
- v_{pmt}: processing time per card of type t on machine m at process p for $p = 1, 2$.
- h_t: oven heat time for type t at process 3.
- $k_{pmtt'}$: setup time when switching from type t to t' on machine (p, m).
- e_w: earliest start time of work order w.
- d_w: due date of work order w.
- π_w: tardiness penalty weight of work order w.
- $a_{tt'} \in \{0, 1\}$: oven compatibility parameter.
- o_m: oven capacity for oven $m \in \mathcal{M}_3$.
- $bc^{\min}_{wt}$, $bc^{\max}_{wt}$: minimum and maximum numbers of active batches for (w, t).
- $bs^{\min}_{wtb}$, $bs^{\max}_{wtb}$: minimum and maximum batch sizes for (w, t, b).
- $N_B := \sum_{w \in \mathcal{W}} \sum_{t \in \mathcal{T}} |\mathcal{B}_{wt}|$: total number of potential batches.
- M: sufficiently large constant.

Decision variables.

- Q_{wtb}: batch size.
- I_{wtb}: equals 1 if batch (w, t, b) is active.
- X_{wtbpm}: assignment variable.
- S_{wtbpm}, C_{wtbpm}: start and completion times.
- R_w: completion time of work order w.
- Z_w: tardiness of work order w.
- $Y_{wtb\,w't'b'\,pm}$: immediate-precedence variable.
- $G_{wtb\,w't'b'\,m}$: oven grouping variable.
- E_{pm}: machine-usage variable.
- F_{wtbpm}, L_{wtbpm}: first and last indicators.
- $D_{wtb\,w't'b'\,m}$: auxiliary oven-ordering variable.

2.3.2 Heuristic Design

An exact mixed-integer linear programming model was first developed to represent the production environment rigorously (see Appendix). The exact formulation determines active batches, machine assignments, sequence decisions, and oven grouping decisions simultaneously. This model is unable to solve instances of real life as they are considerably larger. Therefore, a heuristic solution method was developed for operational use. The heuristic begins by forming feasible batches subject to minimum and maximum tray capacities. Then it constructs an initial sequence using an Apparent Tardiness Cost with Setups logic, which balances urgency and setup effects. The resulting sequence is evaluated by a decoder that simulates the three production stages. At the first two stages, each batch is assigned to the machine that yields the earliest completion time. At the oven stage, compatible batches are grouped into runs while respecting capacity and readiness constraints. The overall sequence is then improved through tabu search, which explores swap-based neighboring solutions. Similar metaheuristic ideas have been shown to be effective in tardiness-oriented scheduling problems with setups (Bilge et al., 2004; Anghinolfi and Paolucci, 2007; Demir et al., 2015).

2.4 Validation of the Proposed Approach

Validation was performed in two stages. First, the internal logic of both the exact model and the heuristic was verified on deliberately constructed small instances. These instances were designed to test precedence relations, batching decisions, machine assignments, setup handling, and oven grouping behavior. In the simplest cases, both approaches produced identical decisions, which confirmed that the heuristic decoder was consistent with the exact formulation on small problems.

Second, the approach was tested using historical company data. The available data set covers completed work orders over a multi-year period and includes order quantities, PCB types, processing-related information, and due dates. Since the observed data do not represent all jobs that actually used the plant capacity during that horizon, the effective capacity of the validation model was reduced in line with expert judgment from the industrial advisor so that the workload level would better reflect reality.

The validation results show that the proposed system behaves consistently with practical scheduling logic. Urgent orders tend to move earlier in the schedule unless setup or penalty trade-offs justify another decision. Jobs with similar setup requirements are often processed consecutively on the same machine when setups are dominant. Likewise, oven grouping decisions do not simply maximize utilization; they also account for batch readiness and delivery urgency. These outcomes indicate that the model captures the main operational trade-offs of the shop floor.

2.5 Integration and Implementation

The decision support system is designed to be integrated with the company's existing planning practice rather than replacing it abruptly. The input and output format is intentionally kept simple so that planners can continue to work with spreadsheet-compatible files while benefiting from algorithmic scheduling support.

In the intended use scenario, planners prepare or export the required order and processing data, run the solution engine, and review the resulting schedule before execution. Since the output includes machine-wise sequences, stage-wise timing information, and work-order completion estimates, planners can interpret the recommendations and revise them when extraordinary situations arise. This improves transparency and facilitates adoption.

The implementation plan also includes a pilot application phase. During this phase, the company can compare historical planning outcomes with schedules suggested by the system for a selected subset of operations. Feedback from planners and the industrial advisor can then be used to refine parameter settings, user-interface components, and rescheduling rules.

2.5.1 Decision Support System Structure

The proposed system is implemented in Python and accepts spreadsheet-based inputs. The input structure includes work orders, quantities, PCB types, batch capacities, processing times, setup times, release dates, due dates, and thermal profile information. The output consists of machine-wise schedules, batch start and completion times, oven-run composition, and work-order level tardiness measures.

For operational use, the system also includes a rescheduling mechanism. At a selected current time, completed work is fixed, ongoing work remains frozen, and only the remaining batches are replanned. This allows the planners to respond to urgent orders or disruptions without destroying the feasibility of operations already in progress.

2.6 Benefits to the Company

The project provides benefits at both operational and managerial levels. Operationally, a better schedule reduces unnecessary setups, improves machine utilization, and lowers the number of late work orders. Managerially, the system provides a repeatable and explainable way of generating schedules, which decreases dependence on purely manual coordination and supports data-driven decision making.

The most visible benefit is the improvement in delivery performance. Reducing weighted tardiness directly supports customer satisfaction and can reduce the risk of contractual penalties in defense-related projects. In addition, by grouping similar jobs more systematically and by using ovens more intelligently, the system increases effective capacity without requiring immediate capital investment.

Table 2.1 summarizes the high-level benchmarking results obtained from the historical comparison.

Table 2.1: Benchmarking summary for the proposed system

Performance measure	Current practice	Proposed system
Total weighted tardiness	204,580	188,067
On-time deliveries	52	57
On-time delivery rate	30.8%	39.0%

Based on these results, the proposed approach yields an 8.6% reduction in total weighted tardiness and an 9.6% increase in the on-time delivery rate. This comparison suggests that the proposed system improves delivery performance while preserving operational realism.

Another important benefit is scalability. The exact model offers a rigorous benchmark for small cases, while the heuristic makes it possible to solve realistic daily instances in acceptable computational times. This balance between rigor and practicality makes the proposed solution suitable for real-world use.

2.7 Conclusions and Future Work

This project addresses a real production planning problem in Karel Elektronik's PCB manufacturing environment by developing a scheduling-oriented decision support system. The study first formalizes the production line as a hybrid flexible flow shop with batching, sequence-dependent setups, and oven compatibility constraints. It then combines an exact model and a practical heuristic to obtain schedules that are both feasible and operationally meaningful.

The results demonstrate that the proposed approach can improve key delivery-related performance measures compared with the current planning practice. Beyond the numerical gains, the project also contributes a structured planning logic that can support planners in a more consistent and transparent way.

Future work may extend the system by incorporating stochastic disruptions, more detailed material-availability constraints, explicit labor limitations, and richer user-interface capabilities. Another natural extension is to integrate the decision support system more closely with ERP data sources so that schedule generation and replanning can be triggered more automatically.

Bibliography

Anghinolfi, D. and M. Paolucci (2007). Parallel machine total tardiness scheduling with a new hybrid metaheuristic approach. *Computers & Operations Research 34*(11), 3471–3490.

Bilge, Ü., F. Kıraç, M. Kurtulan, and P. Pekgün (2004). A tabu search algorithm for parallel machine total tardiness problem. *Computers & Operations Research 31*(3), 397–414.

Demir, H. I., T. Cakar, M. Ipek, O. Uygun, and M. Sari (2015). Process planning and due-date assignment with atc dispatching where earliness, tardiness and due-dates are punished. *Journal of Industrial and Intelligent Information 3*(3), 197–204.

Oujana, S., F. Yalaoui, and L. Amodeo (2019). A linear programming approach for hybrid flexible flow shop with sequence-dependent setup times to minimise total tardiness. *International Journal of Production Research 57*(24), 7651–7667.

Pinedo, M. L. (2016). *Scheduling: Theory, Algorithms, and Systems* (5 ed.). Cham: Springer.

Appendix: Mathematical Model

Objective function.

$$\min \sum_{w\in\mathcal{W}} \pi_w Z_w$$

Constraints.

$$\sum_{b\in\mathcal{B}_{wt}} Q_{wtb} = n_{wt} \qquad \forall w\in\mathcal{W},\ \forall t\in\mathcal{T}$$

$$bs^{\min}_{wtb} I_{wtb} \le Q_{wtb} \le bs^{\max}_{wtb} I_{wtb} \qquad \forall (w,t,b)\in\mathcal{B}$$

$$bc^{\min}_{wt} \le \sum_{b\in\mathcal{B}_{wt}} I_{wtb} \le bc^{\max}_{wt} \qquad \forall w\in\mathcal{W},\ \forall t\in\mathcal{T}$$

$$\sum_{m\in\mathcal{M}_p} X_{wtbpm} = I_{wtb} \qquad \forall (w,t,b)\in\mathcal{B},\ \forall p\in\mathcal{P}$$

$$C_{wtbpm} \le M X_{wtbpm}, \qquad S_{wtbpm} \le M X_{wtbpm} \qquad \forall (w,t,b)\in\mathcal{B},\ \forall p\in\mathcal{P},\ \forall m\in\mathcal{M}_p$$

$$C_{wtbpm} \ge S_{wtbpm} + v_{pmt} Q_{wtb} - M(1 - X_{wtbpm})$$
$$\forall (w,t,b)\in\mathcal{B},\ \forall p\in\{1,2\},\ \forall m\in\mathcal{M}_p$$

$$C_{wtb3m} \ge S_{wtb3m} + h_t - M(1 - X_{wtb3m}) \qquad \forall (w,t,b)\in\mathcal{B},\ \forall m\in\mathcal{M}_3$$

$$S_{wtb1m} \ge e_w - M(1 - X_{wtb1m}) \qquad \forall (w,t,b)\in\mathcal{B},\ \forall m\in\mathcal{M}_1$$

$$S_{wtbpm} \ge C_{wtb(p-1)m'} - M\big(2 - X_{wtbpm} - X_{wtb(p-1)m'}\big)$$
$$\forall (w,t,b)\in\mathcal{B},\ \forall p\in\mathcal{P}\setminus\{1\},\ \forall m\in\mathcal{M}_p,\ \forall m'\in\mathcal{M}_{p-1}.$$

$$R_w \ge C_{wtb3m} \qquad \forall (w,t,b)\in\mathcal{B},\ \forall m\in\mathcal{M}_3$$

$$Z_w \ge R_w - d_w \qquad \forall w\in\mathcal{W}$$

$$\sum_{(w,t,b)\in\mathcal{B}} X_{wtbpm} \le N_B\, E_{pm} \qquad \forall p\in\mathcal{P},\ \forall m\in\mathcal{M}_p$$

$$\sum_{(w,t,b)\in\mathcal{B}} X_{wtbpm} \ge E_{pm} \qquad \forall p\in\mathcal{P},\ \forall m\in\mathcal{M}_p$$

$$S_{wtbpm} \ge C_{w't'b'pm} + k_{pm\,t't} - M(1 - Y_{wtb\,w't'b'\,pm})$$
$$\forall p\in\{1,2\},\ \forall m\in\mathcal{M}_p,\ \forall (w,t,b)\ne(w',t',b').$$

$$Y_{wtb\,w't'b'\,pm} \leq X_{wtbpm} \qquad \forall p \in \mathcal{P},\, \forall m \in \mathcal{M}_p,\, \forall (w,t,b) \neq (w',t',b')$$

$$\sum_{(w',t',b') \neq (w,t,b)} Y_{wtb\,w't'b'\,pm} + F_{wtbpm} = X_{wtbpm} \qquad \forall p \in \{1,2\},\, \forall m \in \mathcal{M}_p,\, \forall (w,t,b)$$

$$\sum_{(w',t',b') \neq (w,t,b)} Y_{w't'b'\,wtb\,pm} + L_{wtbpm} = X_{wtbpm} \qquad \forall p \in \{1,2\},\, \forall m \in \mathcal{M}_p,\, \forall (w,t,b)$$

$$F_{wtbpm} \leq X_{wtbpm}, \qquad L_{wtbpm} \leq X_{wtbpm} \qquad \forall p \in \mathcal{P},\, \forall m \in \mathcal{M}_p,\, \forall (w,t,b)$$

$$\sum_{(w,t,b) \in \mathcal{B}} F_{wtbpm} = E_{pm}, \qquad \sum_{(w,t,b) \in \mathcal{B}} L_{wtbpm} = E_{pm} \qquad \forall p \in \{1,2\},\, \forall m \in \mathcal{M}_p$$

$$S_{wtbpm} \leq S_{w't'b'pm} + M(1 - F_{wtbpm}) + M(1 - X_{w't'b'pm})$$
$$\forall p \in \mathcal{P},\, \forall m \in \mathcal{M}_p,\, \forall (w,t,b) \neq (w',t',b')$$

$$S_{wtbpm} \geq S_{w't'b'pm} - M(1 - L_{wtbpm}) \quad \forall p \in \mathcal{P},\, \forall m \in \mathcal{M}_p,\, \forall (w,t,b) \neq (w',t',b')$$

$$G_{wtb\,w't'b'\,m} \leq a_{tt'} \qquad \forall m \in \mathcal{M}_3,\ \forall (w,t,b) \neq (w',t',b')$$

$$G_{wtb\,w't'b'\,m} \leq \frac{X_{wtb3m} + X_{w't'b'3m}}{2} \qquad \forall m \in \mathcal{M}_3,\ \forall (w,t,b) \neq (w',t',b')$$

$$S_{wtb3m} - S_{w't'b'3m} \leq M\big(1 - G_{wtb\,w't'b'\,m}\big) \qquad \forall m \in \mathcal{M}_3,\ \forall (w,t,b) \neq (w',t',b')$$

$$\sum_{(w',t',b') \neq (w,t,b)} G_{wtb\,w't'b'\,m} \leq o_m - 1 \qquad \forall (w,t,b) \in \mathcal{B},\ \forall m \in \mathcal{M}_3$$

$$\sum_{(w',t',b') \neq (w,t,b)} Y_{wtb\,w't'b'\,3m} + F_{wtb3m} \geq X_{wtb3m} \qquad \forall m \in \mathcal{M}_3,\ \forall (w,t,b)$$

$$\sum_{(w',t',b') \neq (w,t,b)} Y_{wtb\,w't'b'\,3m} \leq o_m - o_m\, F_{wtb3m} \qquad \forall m \in \mathcal{M}_3,\ \forall (w,t,b)$$

$$\sum_{(w',t',b') \neq (w,t,b)} Y_{w't'b'\,wtb\,3m} + L_{wtb3m} \geq X_{wtb3m} \qquad \forall m \in \mathcal{M}_3,\ \forall (w,t,b)$$

$$\sum_{(w',t',b') \neq (w,t,b)} Y_{w't'b'\,wtb\,3m} \leq o_m - o_m\, L_{wtb3m} \qquad \forall m \in \mathcal{M}_3,\ \forall (w,t,b)$$

$$\sum_{(w,t,b) \in \mathcal{B}} F_{wtb3m} \leq o_m, \qquad \sum_{(w,t,b) \in \mathcal{B}} L_{wtb3m} \leq o_m \qquad \forall m \in \mathcal{M}_3$$

$$\sum_{(w,t,b) \in \mathcal{B}} F_{wtb3m} \geq E_{3m}, \qquad \sum_{(w,t,b) \in \mathcal{B}} L_{wtb3m} \geq E_{3m} \qquad \forall m \in \mathcal{M}_3$$

$$S_{wtb3m} \geq C_{w't'b'3m} - M(1 - D_{wtb\,w't'b'\,m}) - M\, G_{wtb\,w't'b'\,m}$$
$$\forall m \in \mathcal{M}_3,\ \forall (w,t,b) \neq (w',t',b')$$

$$S_{w't'b'3m} \geq C_{wtb3m} - M\, D_{wtb\,w't'b'\,m} - M\, G_{wtb\,w't'b'\,m}$$
$$\forall m \in \mathcal{M}_3,\ \forall (w,t,b) \neq (w',t',b')$$

$$Y_{wtb\,w't'b'\,3m} \leq D_{wtb\,w't'b'\,m} \quad \forall m \in \mathcal{M}_3,\ \forall (w,t,b) \neq (w',t',b')$$

$$D_{wtb\,w't'b'\,m} \leq X_{wtb3m} \quad \forall m \in \mathcal{M}_3,\ \forall (w,t,b) \neq (w',t',b')$$

Domains.

$$Q_{wtb},\ S_{wtbpm},\ C_{wtbpm},\ R_w,\ Z_w \geq 0 \qquad \forall (w,t,b) \in \mathcal{B},\ \forall p \in \mathcal{P},\ \forall m \in \mathcal{M}_p$$
$$X_{wtbpm},\ Y_{wtb\,w't'b'\,pm},\ G_{wtb\,w't'b'\,m},\ I_{wtb},\ E_{pm},\ F_{wtbpm},\ L_{wtbpm},\ D_{wtb\,w't'b'\,m} \in \{0,1\}$$

Bireysel Emeklilik Sisteminde Devlet Katkısı Sistem Eniyilemesi | 3

Emeklilik Gözetim Merkezi

Proje Ekibi
İpek Aksoy, Mustafa Gencay Alkan, İrem Atasoy,
Elif Pınar Can, Orbay Ertürk, Ege Güven, Zehra Arzum Küçük

Şirket Danışmanları
Hicran Aktaş Coşar, Barış Çakmak,
Yasemin Yüksel Çuhadar Kütle,
Ege Kaçar ve Samet Özben

Akademik Danışman
Prof. Mustafa Ç. Pınar
Endüstri Mühendisliği Bölümü

Özet

Bireysel Emeklilik Sistemi (BES) kapsamında Devlet, katılımcı katkı paylarının %20'si oranında katkı sağlamakta ve bu tutarlar Devlet katkısı fonlarında yatırıma yönlendirilmektedir. Ancak bu fonlar, katı yatırım kısıtları nedeniyle çoğu zaman negatif reel getiri üretmekte ve performans farklılıkları katılımcılar arasında eşitsizlik yaratmaktadır. Bu çalışma, fon performansı ile yatırım kısıtlarının etkisini inceleyerek en yüksek getirinin kısıtsız senaryoda elde edildiğini göstermektedir. Bu doğrultuda, geçmiş piyasa verileri çerçevesinde getiriyi maksimize eden portföy yapılarının analizine imkan veren bir optimizasyon modeli kullanılmıştır. Model sonuçlarına dayanarak, fonların daha esnek yönetilmesi ve katkıların faizli ve faizsiz olmak üzere iki merkezi fonda toplanması önerilmiştir.

Anahtar Sözcükler: Bireysel Emeklilik Sistemi, Devlet Katkısı, Portföy Optimizasyonu, Değişken Fon, Merkezi Fon, Reel Getiri, Emeklilik Gözetim Merkezi.

System Optimization of State Contribution in the Individual Pension System

Abstract

Under the Individual Pension System (IPS), the government contributes 20% of participants' contributions, and these amounts are invested in state contribution funds. However, due to strict investment constraints, these funds often generate negative real returns, and performance differences create inequality among participants. This study examines the impact of investment constraints on fund performance and shows that the highest returns are achieved under an unconstrained scenario. Accordingly, an optimization model based on historical market data is used to analyze portfolio structures that maximize returns. Based on the model results, it is proposed that funds be managed more flexibly and that contributions be consolidated into two centralized funds, one interest-bearing and one interest-free.

Keywords: Individual Pension System, State Contribution, Portfolio Optimization, Variable Fund, Centralized Fund, Real Return, Pension Monitoring Center.

3.1 Company and System Analysis

3.1.1 Company Information

The Pension Monitoring Center (PMC), established in 2003, is the central institution responsible for operating, monitoring, and developing the IPS in Türkiye, ensuring its transparency, efficiency, and reliability (PMC, 2026a). It oversees State contribution processes, verifies and collects data from pension companies, and ensures the proper execution of fund transfers and allocations. PMC monitors pension fund activities, supports regulatory authorities through reporting and data analysis, and provides audit infrastructure for supervision. It also manages inter-company processes within the system and contributes to policy development and system improvements (PMC, 2026b).

3.1.2 System Analysis

State contribution funds in the Individual Pension System are managed separately by pension companies, although all operate under the same regulatory framework. In this system, the State contributes 20% of the participant's payment, while the annual State contribution amount is limited to 20% of the annual gross minimum wage. These funds are subject to a set of investment constraints, including minimum allocation requirements to government debt and equities, as well as limits on deposits and private debt instruments, as presented in Figure 3.1 (PMC, 2025). In addition, certain asset classes are not permitted under the current regulation, such as foreign assets, precious metals, and derivative instruments. Moreover, when different State contribution funds established under each pension company are managed by different portfolio management companies with distinct strategies, participants may be exposed to varying returns despite being subject to the same system. Within this framework, between 2021 and 2025, the average annual real return of State contribution funds was -6.82% (Turkey Electronic Fund Distribution Platform (TEFAS), 2026).

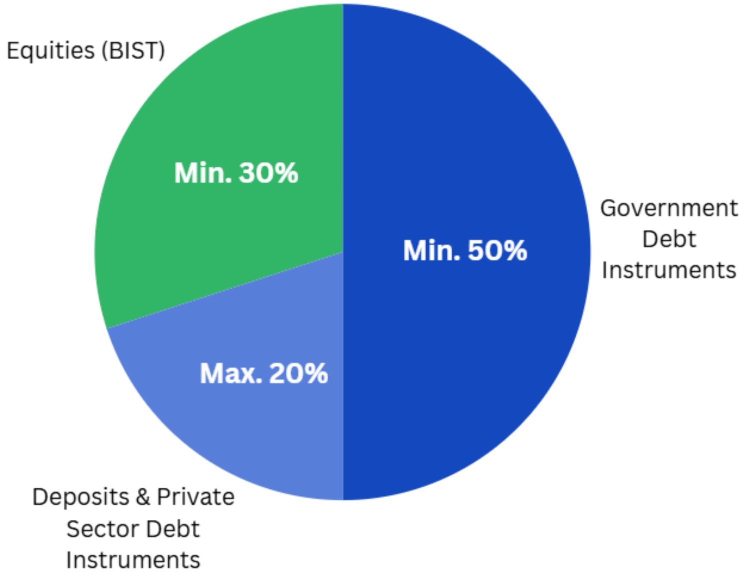

Figure 3.1: Current State Contribution Fund Allocation Constraints

3.2 Problem Definition

The current structure of State contribution funds fails to generate positive real returns, as fund performance remains below inflation due to regulatory and structural limitations. Strict portfolio allocation rules require a large share of assets to be invested in government debt instruments, while limiting diversification and prohibiting alternative investment tools. Between 2013 and 2025, cumulative inflation exceeded 600 percent, while average monthly real returns remained negative at around -0.23% for interest-based funds and -0.48% for non-interest funds. At the same time, the system is fragmented across pension companies. Although all funds operate under identical regulations, differences in asset selection, timing, and operational efficiency lead to unequal participant outcomes despite the same public incentive. Another critical issue is the mismatch between the relatively short maturity of available Treasury instruments, at 3.7 years, and the long-term investment horizon of the IPS (Bilkent University, 2025). These factors weaken accumulation, reduce long-term returns, and contribute to early participant exits. Within this context, system performance is evaluated based on real return, consistency of outcomes across participants, and participant retention. The system also creates the perception of public loss risk, since State contributions may lose real value before being returned to the Treasury in early exit cases (PMC, 2025).

3.2.1 Objectives

The objectives of the proposed framework are to achieve positive real return, ensure capital preservation by reducing potential public loss, establish a fair and standardized return structure by eliminating differences across pension companies, and improve system sustainability by increasing participant retention and supporting the growth of the IPS.

3.2.2 Deliverables

The project has two main deliverables. The first deliverable is the final project report and feasibility analysis. This report includes the objective of the project, key findings, model methodology, and the analyses conducted throughout the study. It also presents the economic impact of the proposed solution, its institutional feasibility, risk and compliance considerations, and a strategic roadmap for implementation. The report provides a structured and comprehensive explanation of the proposed system and its potential benefits.

The second deliverable is a decision support system developed for portfolio optimization to be used by the fund manager, supported by a user interface. The system allows users to run the model for different time periods, select the maximum risk level, and choose the asset classes to be included. Based on these inputs, it produces a portfolio allocation that maximizes real return. In addition, the system provides key outputs such as real return, nominal return, risk degree level, and the contribution of each asset to overall portfolio performance, enabling users to evaluate results and make informed decisions under different scenarios.

3.3 Model and Proposed System

3.3.1 Portfolio Optimization Model

The model is a portfolio optimization framework designed to improve the performance of State contribution funds by maximizing annualized real return. It is formulated using an index-based representation of asset classes and incorporates relevant allocation rules as constraints. While these constraints are strictly enforced in the baseline scenario, sensitivity analyses conducted on them led to their removal in the best-performing scenario to identify more efficient portfolio allocations. Appendix A provides the detailed mathematical formulation of the optimization model, covering volatility, nominal and real return calculations, return annualization, risk degree mapping, as well as the objective function, decision variables, and constraints.

The model follows a structured computational process, as illustrated in Figure 3.2. First, annualized nominal returns are calculated using weekly historical price data. Then, portfolio volatility is estimated to support risk evaluation. Together with annual inflation rates, portfolio constraints, and the selected risk level, these inputs are used in the optimization process, where asset allocations are determined to maximize portfolio real return. The model is implemented in Python as a non-linear optimization problem using the SciPy library and, with its parametric structure, can be used as a decision support tool to evaluate different portfolio scenarios.

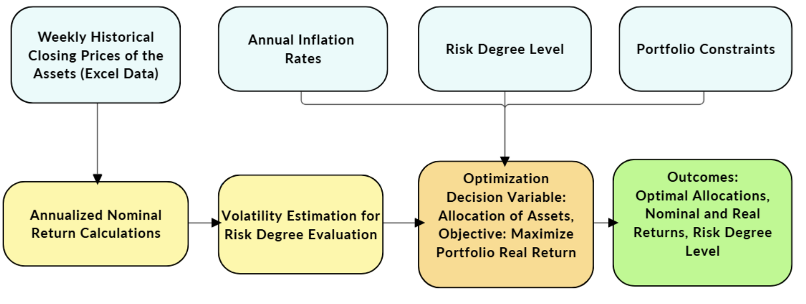

Figure 3.2: Conceptual Model of the Portfolio Optimization Framework

3.3.2 Proposed Solution

Sensitivity analyses on individual portfolio constraints show that these restrictions significantly limit allocation efficiency. When constraints are gradually relaxed, the model consistently shifts toward more flexible allocations that prioritize assets with higher real returns. As a result, the optimized portfolio structure begins to resemble the behavior of variable funds, where asset allocation is not fixed but adjusts according to market conditions. To support this finding, the model results are compared with existing variable pension funds. Data from 26 aggressive variable funds shows an average real return of 14.91%, while the unconstrained model achieves a higher real return of 19.03% (Turkey Electronic Fund Distribution Platform (TEFAS), 2026). In addition, existing funds show significant differences in performance, indicating that outcomes depend largely on fund-level decisions.

These results suggest that the main limitation of the current system is its rigid structure rather than the investment universe itself. Therefore, one of the key proposed solutions is to manage State contribution funds under a variable fund structure, allowing for more flexible and adaptive asset allocation.

Moreover, in the current system, State contribution funds are managed separately by different portfolio management companies, which leads to unequal outcomes for participants receiving the same public incentive. To address this, the proposed solution includes transitioning to a centralized fund structure in which all State contributions are managed under two main funds: one interest-based and one non-interest fund. This structure aims to ensure a more consistent and equitable distribution of returns across all participants.

3.4 Validation

The proposed approach was validated by comparing its results with simulated future outcomes and the actual performance of current State contribution funds.

First, out-of-sample simulation testing was conducted. After the optimized portfolio was developed using historical data, 2,000 future scenarios were generated, and the portfolio's annual nominal return was calculated for each one. Since the in-sample result was not located in the extreme tail of the simulated distribution, it was considered realistic and not overly optimistic.

Second, out-of-sample back-testing was carried out using actual market data from 2024–2025, while the model was trained only on 2021–2023 data. The results were compared with the returns of existing State contribution funds. A very strong positive correlation ($r = 0.89$) was observed, and the t-test showed no statistically significant difference between the model's returns and actual market returns ($p = 0.1226 > 0.05$). This showed that the model produced realistic results and successfully reflected market behavior.

Finally, face validity was supported through expert review. The model structure, assumptions, data, and outputs were evaluated by company experts and the academic advisor, and the approach was found to be consistent with current practice.

Overall, the validation results showed that the proposed approach is reliable and comparable with current operations.

3.5 Implementation and Pilot Study

The proposed approach was designed to fit into the current IPS without changing its main institutional structure. Its main change is replacing multiple State contribution funds managed separately by pension companies with a centralized structure consisting of two funds: one interest-bearing and one interest-free. The overall implementation flow is shown in Figure 3.3. In the proposed system, the funds are legally established by the Republic of Türkiye Ministry of Treasury and Finance, while PMC continues to monitor the system, manage data flows, and support coordination among stakeholders (PMC, 2025).

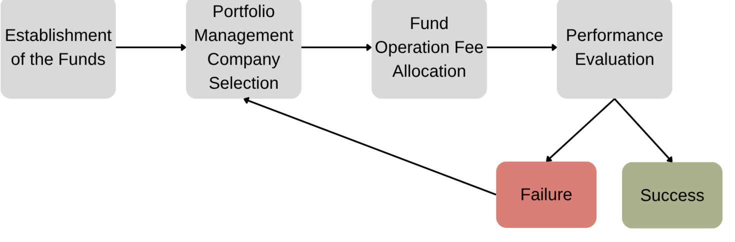

Figure 3.3: Basic implementation flow of the proposed centralized fund structure

In practice, the proposed model preserves the existing roles of pension companies, portfolio management companies, and public authorities. Pension companies continue to collect participant contributions and transfer them into the system, while daily fund management is centralized under a selected portfolio management company. Unlike the current structure, which is subject to strict investment rules, the selected portfolio management company manages these funds under a variable fund structure. This increases asset allocation flexibility, reduces operational complexity, creates a more standardized structure, and helps eliminate return differences caused by company-based fund management. A revenue-sharing mechanism is also proposed to ensure fair income distribution between the managing portfolio management company and pension companies.

The managing institution is proposed to be selected through a transparent evaluation process based on performance, operational capability, and technical adequacy. Following implementation, the selected portfolio management company is proposed to be evaluated annually through a formal performance review mechanism. It is considered unsuccessful if the annual fund return remains below inflation or if its performance falls significantly below the average return of comparable funds. In such cases, the current manager is proposed to be replaced, and a new portfolio management company is selected through the same evaluation process.

In this way, the proposed approach is not only an analytical model but also a practical system design that can be adapted to the company's existing operational framework and support future policy decisions.

To provide initial practical insight into the proposed solution, a pilot study was conducted for the period between January 1, 2026 and March 17, 2026 to examine the nominal returns of the funds. In this pilot application, the State contribution funds were managed under the proposed variable fund structure. During this period, the nominal return of the existing State contribution funds was 7.94%, while the nominal return under the proposed structure reached 20.6% (Turkey Electronic Fund Distribution Platform (TEFAS), 2026). These preliminary results show that managing State contribution funds through a variable fund structure provides stronger return performance compared to the current system. Although the pilot study covers a limited time period, it still provides initial evidence regarding the practical potential of the proposed approach.

3.6 Benefits

The proposed project is expected to provide several benefits to PMC and the overall IPS by improving both fund performance and the system's structure. The main contribution of the proposed model is its effect on return performance. Based on the analysis covering the 2021–2025 period, the average annual real return of the existing State contribution fund structure was -6.82%, whereas the average annual real return of variable funds over the same period was 14.91% (Turkey Electronic Fund Distribution Platform (TEFAS), 2026). This corresponds to a 21.73% improvement in annual real return performance. This result suggests that the proposed structure produces higher returns and supports positive real returns compared with the current system.

In addition to return performance, the centralized structure eliminates the return inequality created by the current fragmented system, in which participants across different pension companies are exposed to different outcomes under the same State contribution mechanism. The introduction of a standardized structure supports fairness and consistency in fund management. In addition, managing the funds under a variable fund structure provides flexibility in asset allocation and alleviates the limitations of the current rigid framework.

Another key benefit of the proposed model is that it eliminates the risk of a public loss perception. Since the current system produces negative real returns, the amounts returned to the Treasury fall below their real value. By delivering stronger and more sustainable real returns, the proposed structure prevents this mismatch and removes the basis for a potential public loss assessment. The proposed revenue-sharing structure also creates a fairer financial framework for pension companies, while the centralized model reduces operational burden by removing the need for each company to manage separate State contribution funds. Overall, the project contributes to PMC by supporting a fairer, more efficient, and more sustainable Individual Pension System.

3.7 Conclusion

This project addressed the main expectations of the organization by identifying the structural weaknesses of the current State contribution system and proposing a practical alternative model. The findings showed that the current structure creates return inequality among participants, limits investment flexibility, and produces weak real return performance. In response, a centralized State contribution fund model managed under a variable fund structure was proposed. The analyses, validation results, pilot study, and implementation planning suggested that this model is both analytically sound and operationally applicable within the existing system framework. In this respect, the project may be considered to have met the organization's expectations by offering a solution that supports fairness, operational simplification, and improved financial outcomes. For future work, the proposed model may be examined further through more detailed regulatory analysis, broader pilot applications, and additional testing under different market conditions, while further studies may also focus on the practical design of the two-fund structure and the long-term effects of the proposed revenue-sharing mechanism.

Bibliography

Bilkent University (2025). 2025-2026 IE 477-478 Öğrencileri İçin Hazırlanan Proje Listesi. https://w3.ie.bilkent.edu.tr/ext-files/project-list-student-version.html. Retrieved October 29, 2025.

PMC (2025). Bireysel Emeklilik Sisteminde Devlet Katkısı Rejimine İlişkin Sistem Analizi ve Tasarımı. Internal company presentation, Emeklilik Gözetim Merkezi. Not publicly available.

PMC (2026a). About PMC. https://www.egm.org.tr/corporate/about-pmc/. Retrieved April 8, 2026.

PMC (2026b). IPS Summary Data. https://www.egm.org.tr/homepage. Retrieved February 5, 2026.

Turkey Electronic Fund Distribution Platform (TEFAS) (2026). Pension funds comparison screen. https://www.tefas.gov.tr/FonKarsilastirma.aspx?type=emk. Retrieved: February 9, 2026.

Appendix: Mathematical Formulation

Sets

- J: Set of all investment assets,

$$\begin{aligned} J = \{&\mathrm{MEVTL}, \mathrm{KDEVL}, \mathrm{ODEGS}, \mathrm{OSABT}, \mathrm{REPBR}, \mathrm{TDTUM}, \mathrm{TTUFE}, \\ &\mathrm{XU100}, \mathrm{XK100}, \mathrm{XBANK}, \mathrm{XUSIN}, \mathrm{EUR}, \mathrm{USD}, \mathrm{GAU\_TRY}, \mathrm{XAGG\_TRY}, \\ &\mathrm{XPD\_TRY}, \mathrm{XPT\_TRY}, \mathrm{EOSTL\_OSBA}, \mathrm{EUSTL\_KAMU}, \mathrm{XGMYO}, \\ &\mathrm{SP\_500}\}. \end{aligned}$$

- $F \subseteq J$: Forbidden assets (not allowed for State contribution funds),

$$\begin{aligned} F = \{&\mathrm{EUR}, \mathrm{USD}, \mathrm{GAU\_TRY}, \mathrm{XAGG\_TRY}, \mathrm{XPD\_TRY}, \mathrm{XPT\_TRY}, \\ &\mathrm{EOSTL\_OSBA}, \mathrm{EUSTL\_KAMU}, \mathrm{XGMYO}, \mathrm{SP\_500}\}. \end{aligned}$$

- $t = 1, 2, \ldots, T$: Weekly observation index.

Parameters

- $P_{j,t}$: Weekly closing value of asset j at week t.
- T: Number of weekly observations (e.g., $T = 260$).
- m: Annualization factor for weekly data ($m = 52$).
- π: Annual inflation rate.
- $RD^{\max}$: Maximum allowed risk degree.
- $\{\tau_1, \ldots, \tau_7\}$: Volatility thresholds for RD classification.

Decision Variables

- w_j: portfolio weight of asset j ($j \in J$).

Return and Risk Computation

Weekly Returns from Closing Prices

$$r_{j,t} = \frac{P_{j,t} - P_{j,t-1}}{P_{j,t-1}}, \qquad \forall j \in J,\ t = 2, \dots, T.$$

Weekly Nominal Portfolio Return

$$r_{p,t} = \sum_{j \in J} w_j \, r_{j,t}, \qquad t = 2, \dots, T.$$

Compounded Total Nominal Return over the Sample
Let $N = T - 1$ be the number of weekly returns.

$$R^{\text{nom}}_{p,\text{tot}} = \prod_{t=2}^{T} (1 + r_{p,t}) - 1.$$

Annualized Nominal Portfolio Return

$$R^{\text{nom}}_{p,\text{ann}} = \left(1 + R^{\text{nom}}_{p,\text{tot}}\right)^{\frac{m}{N}} - 1.$$

Annual Real Portfolio Return

$$R^{\text{real}}_{p,\text{ann}} = \frac{1 + R^{\text{nom}}_{p,\text{ann}}}{1 + \pi} - 1.$$

Annualized Portfolio Volatility

$$\bar{r}_p = \frac{1}{T-1} \sum_{t=2}^{T} r_{p,t}, \qquad \sigma_p = \sqrt{\frac{m}{T-1} \sum_{t=2}^{T} \left(r_{p,t} - \bar{r}_p\right)^2}.$$

Risk Degree Mapping

$$RD_p = k \quad \text{if} \quad \tau_{k-1} < \sigma_p \leq \tau_k, \quad k = 1, \dots, 7, \qquad \tau_0 := 0.$$

To enforce the maximum allowed risk degree: $\sigma_p \leq \tau_{RD^{\max}}$.

Optimization Model

Objective Function

$$\max \quad R^{\text{real}}_{p,\text{ann}}.$$

Constraints

$$
\begin{aligned}
&\sum_{j \in J} w_j = 1, \\
&w_j = 0 \qquad \forall j \in F, \\
&w_j \geq 0 \qquad \forall j \in J, \\
&\sigma_p \leq \tau_{RD^{\max}}, \\
&w_{\text{REPBR}} \leq 0.05, \\
&w_{\text{KDEVL}} + w_{\text{TDTUM}} + w_{\text{TTUFE}} \geq 0.50, \\
&w_{\text{TTUFE}} \leq 0.17\, w_{\text{TDTUM}}, \\
&w_{\text{ODEGS}} + w_{\text{MEVTL}} + w_{\text{OSABT}} \leq 0.20, \\
&w_{\text{XU100}} + w_{\text{XK100}} + w_{\text{XBANK}} + w_{\text{XUSIN}} \geq 0.30.
\end{aligned}
$$

Bireysel Kredilerde Onay Olasılığı Tahminleme Modeli | 4

TeklifimGelsin

Proje Ekibi
Zeynep Aydar, Selin Doğan , Eda Naz Gürel, Eylül Ece Özdemir, Arhan Özturan, Melisa Şahin, Gülce Yıldız

Şirket Danışmanı
Alper Yenigün
Ürün Müdürü

Akademik Danışman
Dr. Öğr. Üyesi Ayşenur Karagöz
Endüstri Mühendisliği Bölümü

Özet

Platformda kullanıcıların kredi başvurularının onaylanma olasılığını tahmin eden bir sistem bulunmamaktadır. Bu durum hem kullanıcı memnuniyetini düşürmekte hem de bankalara yapılan başvurularda verimsizlik yaratmaktadır. Problemi çözmek için, her müşteri ve banka için kredi onay olasılığını tahmin eden bir model geliştirilmiştir. Çalışma kapsamında farklı makine öğrenmesi yöntemleri denenmiş ve karşılaştırılmıştır. Son aşamada, karar ağacının ortaya çıkardığı değişken etkileşimleri lojistik regresyon ile birleştirilerek hibrit bir model kurulmuştur. Böylece model anlaşılır hale gelmiş daha iyi sonuçlar vermiştir. Model gerçek verilerle test edilmiş ve mevcut sistemle karşılaştırılmıştır. Model, kullanıcıları daha uygun seçeneklere yönlendirecek ve onların başvuru süreçlerinde daha bilinçli kararlar almalarını destekleyecektir.

Anahtar Sözcükler: Kredi skoru, kredi onay tahmini, finansal veri analizi, makine öğrenmesi algoritmaları, hibrit lojistik regresyon modeli.

Loan Approval Probability Prediction Model for Retail Loans

Abstract

TeklifimGelsin currently lacks a predictive system to evaluate the loan approval probability of its users. This leads to customer dissatisfaction and inefficiencies in the bank application processes. To address this problem, we developed a predictive model that estimates the probability of loan approval for each customer-bank pair. Multiple machine learningmodels were implemented and compared. In the final approach, a hybrid logistic regression model was developed by combining logistic regression with interaction effects identified through decision tree analysis, improving both interpretability and predictive performance. The model was tested using real data and compared with the existing system. The results show that the model can help guide users toward more suitable options and support more informed decision-making during the application process.

Keywords: Credit score, loan approval prediction, financial data analysis, machine learning algorithms, hybrid logistic regression model.

4.1 About the Company

TeklifimGelsin is a Turkey based financial service provider founded in 2020 that delivers personalized banking offers through digital platforms. The company connects individual customers with banks, enabling users to compare various financial products such as loans, deposit accounts, credit cards, and business support options. It also provides investment calculation services in collaboration with a digital investment platform.

TeklifimGelsin collaborates with multiple banks and financial institutions to create a broad marketplace of financial offers, allowing users to evaluate different options and select the most suitable one for their needs.

4.2 System Analysis

Currently, users can browse loan offers on TeklifimGelsin and apply to them through the platform. TeklifimGelsin has an ideal process for users who want to apply for financial products. In this process, users first create an account and generate a credit report by sharing their personal and financial information. This data is used to obtain the user's Findeks score and is combined with the platform's own analysis to evaluate their financial situation. The result is shown as a letter grade from A (best) to E (worst), and it is further supported with + or - signs based on the platform's own insights. The report also gives suggestions on how users can improve their financial status.

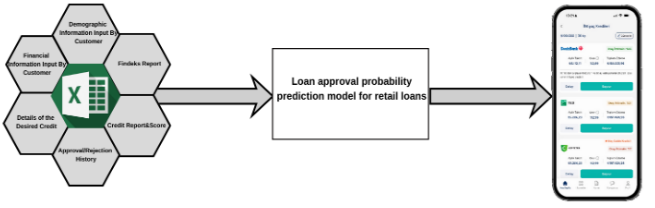

Figure 4.1: Flow Chart of the System

In the real world, many users deviate from the ideal path. Some apply without generating a credit report, while others ignore the information and apply regardless of their financial status. As a consequence of this behavior, users' satisfaction with TeklifimGelsin is affected when their applications are rejected, even though the banks make the final decision.

A critical observation from analyzing the platform's data was that the rejection-to-approval ratio in the historical dataset stands at approximately 12.2 : 1, which indicates that the majority of applications submitted through the platform are ultimately rejected. This ratio shows the scale of the mismatch between users and the loan products they apply for.

4.3 Problem Definition

After analyzing the system, the main problem TeklifimGelsin faces is identified as the lack of guidance in matching the right customer with the right bank. Although many options are provided, the platform does not have a systematic way to evaluate whether a user will be approved for a specific loan. This is mainly because the decision mechanisms of banks are not known, and the importance of different customer attributes cannot be clearly determined. TeklifimGelsin tries to address this through credit reports based on Findeks scores, but these reports are not enough. As a result, users apply for loans they are unlikely to be approved for, leading to high rejection rates, lower user satisfaction, and inefficiencies for banks. Therefore, the main goal of this project is to reduce customer dissatisfaction by improving customer bank matching.

4.4 Proposed Solution Strategy

4.4.1 Critical Assumptions

Six main assumptions are defined to guide the solution. First, it is assumed that user data is complete and accurate. Second, the provided dataset is assumed to represent the target population. Third, indirect applications are excluded from evaluation. Fourth, it is assumed that users correctly understand the chances of getting approval. Fifth, a user-specific acceptance threshold μ is defined, which means that each customer only applies if the approval probability is higher than their own willingness threshold. Sixth, a standard decision threshold τ is used for model classification and performance evaluation. Applications with a predicted approval probability greater than or equal to $\tau = 0.5$ are treated as accepted; otherwise, they are rejected.

4.4.2 Constraints and Objective

The solution is developed under several constraints. The first constraint is that TeklifimGelsin has no control over the bank or product selected by the customer. The second is that the decision mechanisms and customer data used by banks are not fully known, which limits the ability to model approval decisions accurately.

The main objective of the project is to maximize the probability of correct predictions by improving the accuracyof loan approval estimations.

4.4.3 Solution Approach

Conceptual Model

The solution takes each user's personal profile, financial history, credit report data, desired loan amount, and loan characteristics (maturity, interest rate, bank) as input and calculates an estimated approval probability for each available bank offer.

At the beginning, we tried four different machine learningmodels, which are: Logistic Regression, Decision Tree, XG Boost, and Random Forest. After consultation with academic and industrial advisors, a combined approach was chosen as the core modeling strategy. Rather than utilizing a single machine learning algorithm, the hybrid logistic regression model combines the strengths of two methods: clean visualization of significant variables and their interactions through decision tree and easy interpretability of the logistic regression (Yang, 2024; Lotfi, 2024). In this approach, a hybrid logistic regression model combines the interactions found by the decision tree with a logistic regression model. This lets the model capture variable interactions while still being easy to understand, and gives approval probability estimates for each customer bank pair.

Mathematical Model

Five machine learning models were implemented and compared on a dataset of 55,120 loan application records (38 features) provided by TeklifimGelsin, using a 70–30 train–test split. Each application is represented as a pair $(x^{(i)}, y^{(i)})$, where $x^{(i)}$ is the feature vector of customer and offer attributes, and $y^{(i)} \in \{0, 1\}$ indicates rejection or approval.

Logistic regression model estimates the approval probability by applying the sigmoid function to a linear combination of input features:

$$\hat{p} = \frac{1}{1 + \exp\left(-\left(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_n x_n\right)\right)} \tag{4.1}$$

Each coefficient β_j captures the direction and magnitude of a feature's effect on the log-odds of approval, making the model highly interpretable. Parameters are estimated by minimizing the binary cross-entropy loss over the training set.

The chosen model is a hybrid logistic regression model that proceeds in two stages. First, a pruned decision tree is fitted to the training data. By analyzing consecutive splits along the same branch of the tree where the same pair of variables jointly drives classification outcomes, significant interaction terms are identified. ANOVA tests are conducted to confirm that each candidate interaction term provides a statistically significant improvement (p-value < 0.05) before it is retained.

In the second stage, these interaction terms are added to the logistic regression modelas additional engineered features. The resulting model is:

$$\hat{p} = \frac{1}{1 + \exp\left(-\left(\beta_0 + \sum_j \beta_j x_j + \sum_{(j,k)\in\mathcal{I}} \gamma_{jk} x_j x_k\right)\right)} \tag{4.2}$$

where $\mathcal{I}$ denotes the set of significant interaction terms identified by the decision tree, and γ_{jk} are the corresponding interaction coefficients estimated during logistic regression fitting.

As discussed, a severe class imbalance was present in the dataset, with rejected applications outnumbering approvals at a ratio of 12.2:1. To prevent the model from defaulting to predicting rejection for all inputs, the Synthetic Minority Oversampling Technique (SMOTE) is applied during training, generating synthetic approved-application samples to balance the class distribution. A classification threshold of $\tau = 0.5$ is used to convert predicted probabilities into binary decisions.

Seasonality Analysis

To assess the robustness of the model across different time periods, a seasonality analysis was conducted on the historical data. The results revealed that approval rates are higher at the beginning of the year and decrease in the middle. At the bank level, different seasonal patterns were observed, and banks were classified based on these patterns. To capture time-varying effects, a rolling logistic regression model with a sliding window approach was applied. However, the results were not statistically significant, meaning that seasonalitydoes not have a strong impact on the model. Therefore, the proposed hybrid logistic regression model is considered to be robust to seasonal variations.

4.5 Validation

The model was validated to ensure the consistency between its predictions and real outcomes from the historical data.

1. **Data Splitting:** We split the data given to us by TeklifimGelsin by 70% to 30% for training and testing the model. This process ensured that the model was trained on a large sample of the dataset, whereas the remaining portion was independently used for testing to understand how well the model works in the unseen part of the data.

2. **Confusion Matrix:** We used a confusion matrix to look at the correct and incorrect predictions and find the false positives and false negatives. It was also used to check the performance metrics.

3. **Accuracy, Precision, Recall, F1-score Evaluation:**Performance was evaluated based on test set results, as large differences between training and test performance may indicate overfitting. When observed, these differences were analyzed in terms of class imbalance or data noise.

Model	Accuracy	Precision	Recall	F1- Score
Decision Tree	80.10%	28.90%	82.30%	0.428
XG Boost	81.30%	26.10%	72.00%	0.383
Random Forest	77.40%	22.60%	74.10%	0.346
Logistic Regression	84.70%	18.50%	20.50%	0.195
Hybrid Model	76.90%	26.70%	89.60%	0.411

Figure 4.2: Comparison of Model Performance Metrics

4. **Cross Validation:** We performed 10-fold cross-validation so that we can evaluate the model's performance and ensure it is stable. The model was trained and tested on different subsets of the data, resulting in an average accuracy of 90.95%

5. **Expert Validation:** We also validated the model through discussions with our academic advisor and the company. Based on this feedback, the hybrid logistic regression model was selected as the suitable approach.

4.6 Benchmarking

First, to evaluate the impact of the model, we tested it on historical data. Regarding the company's expectations, we worked with the users whose approval probabilities are higher than $\mu = 0.5$. Consequently, disencouraging the users with approval probabilities less than $\mu = 0.5$ increased the approval rate from 9% to 27.6% compared to the current system.

Secondly, the model was compared with the company's internal prediction system. The results showed that the model significantly improves key performance metrics. In particular, recall increased from 28.8% to 61.5%. This means that the model is significantly better at identifying approved applications. In addition, the true negative rate also increased substantially, from 45.7% to 72.4%, showing that the model is more successful in correctly identifying rejected applications. These improvements show that the model increased the correct prediction of approved and rejected applications. Such a significant improvement stems from the increased complexity of our model compared to the company's internal system in terms of the variables considered.

4.7 Pilot Study

Following the model finalization, a pilot study was conducted in collaboration with TeklifimGelsin to evaluate real-world performance on live platform data. The model was applied internally by the company to a selected group of customers. Users who were identified by the model to have high approval probabilities for specific loans were notified via push notifications to encourage them to apply. The outcomes of these encouraged applications were then tracked and assessed to determine whether the model's high-probability predictions result in real-world approvals.

During the pilot, the rejection-to-approval ratio was also monitored. The implementation results aligned with the expectation of a decrease in the ratio: the observed rejection-to-approval ratio dropped to approximately

7.3:1, compared to the baseline of 12.2:1 in the historical data. This result validated the model's effectiveness in the real world.

4.8 Benefits to the Company

By offering loan options to customers based on predicted approval probabilities and directing them toward loans with higher chances of approval, the proposed system increases user satisfaction. Users focus on choices that are a better fit for their profile instead of applying for loans at random and getting rejected. Furthermore, by matching the right customer with the right bank, the strength of the relationship between TeklifimGelsin and partner banks may increase.

4.9 Conclusion

The project met the expectations of TeklifimGelsin and provided additional insights into the approval likelihood of the applicant. A machine learning-based probability prediction model, the hybrid logistic regression model, increased the company's success in identifying approved applications by 32.7% and rejected applications by 26.7%. Furthermore, the model lowered the observed rejection-to-approval ratio from 12.2:1 to approximately 7.3:1.

The model is long-lasting, understandable, and easily adaptable with minimal necessary data cleaning to the given dataset. Future work may include training the model on years-long data for further seasonality analysis and omitting the proposals from the page for applicants with very low approval probability from specific banks to enhance the model's effectiveness.

Bibliography

Lotfi, E. (2024). Predicting credit card approval using machine learning techniques. *1*, 1–13.

Yang, C. (2024). Research on loan approval and credit risk based on the comparison of machine learning models. In *SHS Web of Conferences*, Volume 181, pp. 02003. EDP Sciences.

Appendices

4.A Logistic Regression Model

Logistic Regression Model

Notation

- $(x^{(i)}, y^{(i)})$: Observation i, for $i = 1, \dots, m$

- $x^{(i)} = (x_1^{(i)}, \ldots, x_n^{(i)})$: Feature vector of observation i
- $y^{(i)} \in \{0, 1\}$: Binary outcome (1: approval, 0: rejection)
- $\beta = (\beta_0, \beta_1, \ldots, \beta_n)$: Model coefficients
- $z^{(i)}$: Linear predictor for observation i
- $\hat{p}^{(i)}$: Predicted probability of approval
- $\ell^{(i)}(\beta)$: Loss of observation i
- $J(\beta)$: Average loss

Model Definition:

$$z^{(i)} = \beta_0 + \sum_{j=1}^{n} \beta_j x_j^{(i)} \quad \text{and} \quad \hat{p}^{(i)} = \frac{1}{1 + e^{-z^{(i)}}} \qquad \forall i = 1, \ldots, m$$

Objective Function:

$$\min_{\beta} \quad J(\beta) = -\frac{1}{m} \sum_{i=1}^{m} \left[y^{(i)} \log \hat{p}^{(i)} + (1 - y^{(i)}) \log(1 - \hat{p}^{(i)}) \right]$$

Constraints:

$$\beta_j \in \mathbb{R} \qquad \forall j = 0, 1, \ldots, n$$

4.B Top 5 Most Important Variables

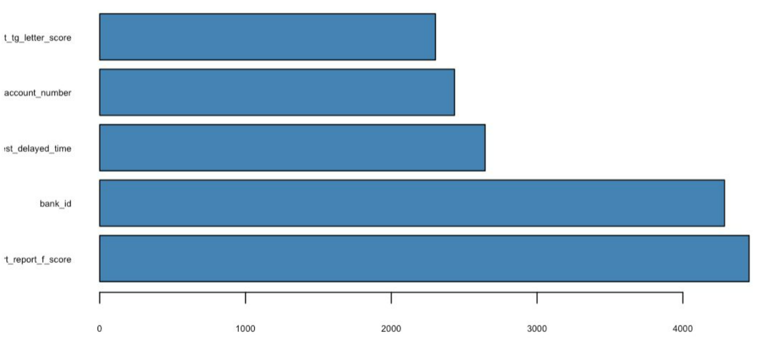

5 | Depo İşgücü Kaynak Planlama

Roketsan

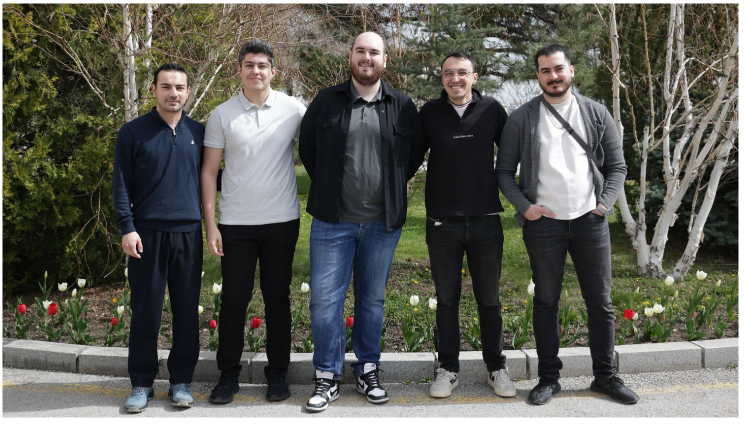

Proje Ekibi
Mert Coşkun, Abdurrahman Çakılca, Birkan Karakaş, Ali Kastal
Hüseyin Kaan Kılıç, İhsan Öğ, Toprak Tuntaş

Şirket Danışmanı
Tevfik Buğra Ünsal
Ostim Depo Operasyonları Birimi
Mühendis

Akademik Danışman
Doç. Dr. Emre Nadar
Endüstri Mühendisliği Bölümü

Özet

Bu çalışma, ürün çeşitliliği ve talep dalgalanmalarının aşırı fazla mesai ve iş yükü dengesizliklerine yol açtığı Roketsan merkez ambarındaki iş gücü planlama verimsizliklerini ele almaktadır. Çalışma kapsamında; geçmiş iş emirlerini zaman etütleri vasıtasıyla işgücü-saat verisine dönüştüren, iş yükü öngörüsü için Winter-Holt üçlü üssel düzleştirme yöntemini entegre eden ve dinamik iş gücü çizelgeleme için bir tamsayılı programlama modeli kullanan bir karar destek sistemi önerilmektedir. Geliştirilen çözümün, fazla mesaiyi ve iş yükü dengesizliklerini azaltmada etkin bir araç olduğu gösterilmektedir.

Anahtar Sözcükler: İşgücü planlama, Winter-Holt tahminleme, tamsayı programlama.

Warehouse Workforce Resource Planning

Abstract

This study addresses workforce planning inefficiencies at Roketsan's central warehouse, where product variety and demand volatility create excessive overtime and workload imbalances. We propose a decision support system (DSS) that translates historical work orders into labor-hours via time studies, integrates Winter-Holt's triple exponential smoothing for workload forecasting, and utilizes an integer programming (IP) model for dynamic workforce scheduling. The proposed framework minimizes total overtime hours and ensures equitable work load distribution while strictly adhering to Turkish Labor Law and operational constraints.

Keywords: Workforce planning, Winter-Holt's forecasting, integer programming.

5.1 Company Profile and Problem Definition

5.1.1 Company description

Established in 1988, Roketsan is a prominent international defense corporation and a cornerstone of Turkey's military industry, specializing in rocket, missile, and precision strike systems. To maintain production continuity across its manufacturing systems, Roketsan relies heavily on its Ostim Central Warehouse. This 10,000-square-meter facility serves as a critical logistical hub, receiving and storing raw materials, semi-finished goods, and spare parts from numerous local subcontractors before they are dispatched to the main production plants. The warehouse employs a total of 25 blue-collar warehouse workers who manage daily operations including material receiving, addressing, and order preparation.

5.1.2 System analysis and problem definition

The warehouse currently operates under highly volatile demand conditions. Order inflows are often irregular and project-based, lacking prior information regarding the physical characteristics or volume of the incoming materials. Additionally, there is a tendency to keep excess inventory in the warehouse, which leads to high occupancy rates and space limitations. These conditions create severe bottlenecks in material movement and make it nearly impossible for personnel to plan resource allocation in advance. Consequently, the facility struggles to complete its daily workload within regular working hours, leading to excessive and inequitable overtime among the workforce.

5.1.3 Deliverables and performance measures

To address this workforce planning inefficiency, this project develops a data-driven DSS. The system utilizes a Winter-Holt time-series model to forecast the upcoming weekly workload in labor-hours and integrates an IP model to automatically optimize daily shift schedules. The primary performance measures for this project are minimizing total overtime hours and reducing workload imbalance (variance in total weekly working hours) among employees.

5.2 Solution System Design and Modeling

The proposed solution for Roketsan's workforce planning challenges is a fast, data-driven, and optimization-based DSS. The conceptual workflow begins with the ingestion of six months of historical operational data, where past work orders are translated into labor-hours by correlating task volumes with standard times derived from on site time studies. This transformation is essential for shifting from qualitative descriptions to the quantitative modeling required for effective industrial engineering applications.

5.2.1 Forecasting model

To estimate the upcoming week's daily workload (C_k), the system utilizes the Winter-Holt's Triple Exponential Smoothing method, motivated by its proven ability to align resource capacity with fluctuating demand in industrial environments. Specifically, the model's additive seasonality and damped trend configurations were selected following the methodology of Nahmias and Olsen (2015) to decompose operational data into level, trend, and seasonal components, ensuring that workload predictions are mathematically grounded before being applied to resource allocation.

- y_t: observed daily workload at day t,
- l_t: level component,
- b_t: trend component
- s_t: seasonal component,
- $m = 7$: seasonal period of one week.
- α, β, γ: smoothing parameters for level, trend, and seasonality
- ϕ: damping parameter ($0 < \phi < 1$), which controls the decay of the trend

$$l_t = \alpha(y_t - s_{t-m}) + (1-\alpha)(l_{t-1} + \phi b_{t-1})$$

$$b_t = \beta(l_t - l_{t-1}) + (1-\beta)\phi b_{t-1}$$

$$s_t = \gamma(y_t - l_{t-1} - \phi b_{t-1}) + (1-\gamma)s_{t-m}$$

The resulting h-step-ahead forecast is calculated by:

$$\hat{y}_{t+h} = l_t + \left(\sum_{i=1}^{h} \phi^i\right) b_t + s_{t-m+h}$$

The output of the forecasting model is a vector of six predicted daily workload values corresponding to the next operational week (Monday–Saturday). This C_k becomes the primary input for the workforce optimization model described in the next subsection.

5.2.2 Optimization model

The core of the system is an IP model designed to minimize total overtime hours for the 25-person workforce while satisfying the forecasted daily workload. The choice of an IP framework is supported by established literature, such as Winston (2004), which provides the foundation for formulating objective functions and constraints in workforce and capacity planning problems. The model assigns workers to discrete overtime tiers through binary decision variables, ensuring that total available labor-hours meet or exceed the daily forecasted demand. A critical feature of this approach is the inclusion of a fairness parameter (Δ), which limits the maximum difference in total weekly working hours between any two employees to promote organizational justice and prevent fatigue. This balancing mechanism ensures that the optimization results are not only efficient but also socially sustainable. Further the model emphasizes compliance with regulations set by the Republic of Turkey, Ministry of Justice (2003) via the overtime limiting α parameter.

Model parameters and decision variables:

- $i = 1, 2, \ldots, 25$: Index representing each worker, where there are 25 employees in total.

- $k = 1, 2, 3, 4, 5, 6$: Index representing the days of the week, corresponding to Monday through Saturday.

- $X_{i,k} = \begin{cases} 1, & \text{if worker } i \text{ works a regular shift on day } k \\ 0, & \text{otherwise} \end{cases}$

- $Y_{1,i,k} = \begin{cases} 1, & \text{if worker } i \text{ works 1 hour overtime on day } k \\ 0, & \text{otherwise} \end{cases}$

- $Y_{2,i,k} = \begin{cases} 1, & \text{if worker } i \text{ works 2 hour overtime on day } k \\ 0, & \text{otherwise} \end{cases}$

- $Y_{3,i,k} = \begin{cases} 1, & \text{if worker } i \text{ works 3 hour overtime on day } k \\ 0, & \text{otherwise} \end{cases}$

- $Y_{i,k} = \begin{cases} 1, & \text{if worker } i \text{ works any amount of hours overtime on day } k \\ 0, & \text{otherwise} \end{cases}$

- C_k: The total workload (in labor-hours) required to be completed on day k.

- Z_i: Total weekly working hours (regular + overtime) for worker i.

- Δ: Maximum allowed difference between total weekly hours of any two workers.

- α: Maximum allowed weekly overtime hours per worker.

Objective function: The objective function minimizes the total amount of overtime across all six working days.

$$\min \sum_{i=1}^{25} \sum_{k=1}^{6} Y_{1,i,k} + 2 \sum_{i=1}^{25} \sum_{k=1}^{6} Y_{2,i,k} + 3 \sum_{i=1}^{25} \sum_{k=1}^{6} Y_{3,i,k}$$

Constraints:

$$7.5 \sum_{i=1}^{25} X_{i,k} + \sum_{i=1}^{25} Y_{1,i,k} + 2 \sum_{i=1}^{25} Y_{2,i,k} + 3 \sum_{i=1}^{25} Y_{3,i,k} \geq C_k, \quad \forall k = 1, 2, \ldots, 6$$

$$Y_{1,i,k} + Y_{2,i,k} + Y_{3,i,k} = Y_{i,k}, \quad \forall i = 1, 2, \ldots, 25 \; \forall k = 1, 2, \ldots, 6$$

$$Y_{i,k} \leq 1, \quad \forall i = 1, 2, \ldots, 25; \; \forall k = 1, 2, \ldots, 6$$

$$X_{i,k} \geq Y_{i,k}, \quad \forall i = 1, 2, \ldots, 25; \; \forall k = 1, 2, \ldots, 6$$

$$7.5\sum_{k=1}^{6}X_{i,k}+\sum_{k=1}^{6}Y_{1,i,k}+2\sum_{k=1}^{6}Y_{2,i,k}+3\sum_{k=1}^{6}Y_{3,i,k}=Z_i,\quad \forall i=1,2,\ldots,25$$

$$Z_i-Z_j\leq\Delta,\quad \forall i,j=1,2,\ldots,25;$$

$$Z_i\leq 7.5\sum_{k=1}^{6}X_{i,k}+\alpha,\quad \forall i,=1,2,\ldots,25;$$

$$X_{i,k},Y_{1,i,k},Y_{2,i,k},Y_{3,i,k}\in\{0,1\},\quad \forall i=1,2,\ldots,25;\ \forall k=1,2,\ldots,6$$

5.2.3 Software environment

Regarding the technical implementation; the development, and solution of these models required a multi-platform strategy tailored to the company's operational environment. While the initial data analysis, parameter fine-tuning, and forecasting logic were developed using Python with libraries such as statsmodels and PuLP+CBC, the final deliverable was transitioned to a Microsoft Excel-based interface. This choice was motivated by the company's IT security policies, which restrict external software installations. Consequently, the decision support tool utilizes VBA for automated data processing and OpenSolver as the primary optimization engine to solve the IP model. This integrated platform ensures that warehouse managers can manage data updates and generate feasible schedules within a familiar and secure software environment, facilitating the seamless integration of the DSS into Roketsan's existing infrastructure.

5.3 System Validation and Benchmarking

To ensure the credibility and feasibility of our proposed DSS, we conducted a broad validation process. This phase measured how our data-driven approach compares against Roketsan Ostim Warehouse's manual planning, proving the model's logic and predictive accuracy under real-world conditions.

5.3.1 Comparative analysis

Current scheduling relies on a manual, heuristic approach where overtime is assigned without strict mathematical optimization procedures. Due to volatile, project-based material flows, this method struggles to balance workloads, leading to employee imbalances and excessive overtime.

Our approach replaces experience-based decision-making with an IP optimization model. To prove its superiority, we used historical work orders provided by our Industrial Advisor (IA), Mr. Tevfik Buğra Ünsal. The DSS successfully converted raw transaction data into labor-hours, generating an optimized, compliant workforce schedule without failures. Stress tests confirmed that the system correctly identifies and prevents constraint violations. By using real data, we demonstrated that our model eliminates manual errors, adheres to Turkish Labor Law, and guarantees fair task distribution.

5.3.2 Validation of workload calculation

To further validate our data processing logic, we compared calculated historical workloads against actual recorded overtime hours. By multiplying historical work orders by our established time study standards, we derived the total workload for past periods. Our findings showed that the calculated workload was marginally lower than the actual overtime hours utilized. This small, consistent variance confirms that our methodology for calculating labor requirements through time studies is accurate and reflects the true operational demands of the warehouse.

5.3.3 Forecasting model validation

A core advantage of our system is anticipating workload rather than reacting to it. To validate our Winter-Holt's Exponential Smoothing engine, we employed a back-testing methodology. We split the historical dataset into two: 24 weeks for training and 1 for testing. Comparing predicted labor-hours against actual workloads yielded a Mean Absolute Percentage Error (MAPE) of 19.7%. In forecasting a MAPE below 20% is generally considered successful. This validates that our forecasting module provides a reliable, mathematically grounded input.

5.3.4 Optimization model robustness and face validity

Face Validity sessions with Mr. Ünsal confirmed that the model's constraints accurately mirror actual operational boundaries and legal obligations. Furthermore, we subjected the IP model to various tests to prove its mathematical and logical validity.

5.3.5 Scalability and future adaptability

The system's success is not merely a curve-fit; it is highly generalizable. The forecasting engine uses a "rolling window" approach, automatically adapting to new trends as weekly data is updated. Additionally, parameters like workforce size and overtime limits are adjustable via an Excel interface, allowing the DSS to adapt to warehouse expansions or policy changes. Consequently, the benefits of minimized overtime, legal compliance, and workload equity will translate to Roketsan's daily operations for years to come.

5.4 Integration and Implementation

5.4.1 Integration strategy

The integration of the DSS into Roketsan's infrastructure is designed to be non-intrusive and fully compliant with the strict security protocols and IT policies of the defense industry.

- **Excel based environment:** To bypass restrictions on external software installations, the system is implemented entirely within Microsoft Excel using VBA. This ensures that authorized personnel can execute the models on corporate workstations without requiring administrative privileges or external library dependencies.

- **Optimization via OpenSolver:** The IP model utilizes Excel Solver structures and is solved using the OpenSolver add-in. This configuration was confirmed by the company to be compatible with their existing IT infrastructure.

- **Standardized data interface:** A dedicated Excel template facilitates the transfer of historical work orders and time study data from the company's current systems.

- **Automated workflow:** The DSS is programmed to automatically read the standardized input files, process historical data into labor-hour requirements, and generate the necessary C_k workload inputs to produce the final weekly overtime schedule.

This streamlined approach transforms the scheduling process into an automated, data-driven operation that remains adaptable to future changes in workforce size or warehouse layout.

5.4.2 Pilot study: shadow implementation phase

A one-week "Shadow Implementation" (Parallel Run) has been conducted to validate the performance of the VBA-based DSS in a live environment without disrupting daily operations. This phase is designed to ensure the system is technically functional and capable of handling live data streams from the warehouse. During this period, the DSS operated in the background, ingesting real-time work orders and generating optimization schedules while the warehouse continues to be managed manually. This methodology allows for a direct quantitative comparison of "Total Overtime Hours Assigned" and "Worker Load Balance" between the manual and optimized systems. Following the pilot, a review meeting was held with the IA to refine the model's constraints based on these real-world results.

5.5 Benefits to the Company

The implementation of the proposed workforce optimization system transitions the warehouse from a manual, heuristic-based scheduling approach to an automated, data-driven operation, offering several key benefits to Roketsan:

- **Reduction in total overtime (Cost efficiency):** By rigorously forecasting the daily workload and optimally allocating shifts, the system minimizes unnecessary overtime hours. This direct reduction in excess labor hours is expected to lower operational costs. Retrospective analysis using historical data provided by the IA indicates that the proposed methodology could achieve an 11% reduction in total overtime hours compared to existing manual practices within the analyzed period.

- **Workload equity and employee satisfaction:** The system incorporates a balancing parameter ($\Delta = 5$ hours) to ensure a fair distribution of total weekly working hours among the 25 warehouse workers. This constraint addresses a significant inequity in the current manual scheduling process, where the average weekly discrepancy between the most and least overtime-assigned employees is observed to be 25 hours (80% decrease). By drastically reducing this variance, the equitable task distribution significantly enhances employee morale, prevents worker fatigue, and promotes organizational justice.

- **Increased operational focus:** Automating the complex task of shift planning frees up valuable time for warehouse management and engineers. The transition from manual scheduling -which previously required several hours of administrative effort each week- to the automated DSS allows optimal schedules to be generated within minutes. This shift enables personnel to focus on continuous process improvements and strategic initiatives rather than repetitive scheduling overhead.

5.6 Conclusion

This project has successfully developed a data-driven DSS to address the workforce planning and workload distribution challenges at Roketsan's central warehouse. By transitioning from a manual scheduling approach to an IP-based framework, the project satisfies the organization's primary expectations of predictability and improvement in working conditions.

The integration of Winter-Holt's triple exponential smoothing for forecasting and the IP model for scheduling ensures that resource allocation is both mathematically grounded and legally compliant. The proposed system directly addresses the identified problem of excessive overtime by minimizing unnecessary labor hours and promoting organizational justice through the equitable distribution of tasks among the 25-person workforce. Furthermore, by implementing the solution within a secure, Excel-based environment using VBA and OpenSolver, the project adheres to the company's strict IT security protocols while providing a user-friendly interface for warehouse managers.

Bibliography

Nahmias, S. and T. L. Olsen (2015). *Production and operations analysis: Strategy, quality, analytics, application* (7 ed.). Waveland Press.

Republic of Turkey, Ministry of Justice (2003). Labor law no. 4857.

Winston, W. L. (2004). *Operations research: Applications and algorithms.* Duxbury Press.

6 Fabrika 6 Malzeme Dağıtım Ağı Eniyilemesi

Beko Buzdolabı İşletmesi

Proje Ekibi
Buket Ay, Berke Ekinci, Yiğit Eyi, Çağan Berkin Öztürk,
Tuna Pulak, Musa Çağan Terzioğlu, Selin Dila Uslu

Şirket Danışmanı
Hüseyin Çoker
Kıdemli Uzman Mühendis

Akademik Danışman
Doç. Dr. Firdevs Ulus
Endüstri Mühendisliği Bölümü

Özet

Bu proje, Beko Buzdolabı İşletmesi Fabrika 6'daki iç lojistik sisteminde, gerçek malzeme akış desenlerine dayalı otomatik yönlendirmeli araç hat yönlerinin eniyilenmesine odaklanarak kritik bir verimsizliği ele almaktadır. Sorunun çözümü için problem, bir "Ağırlıklı Çizge Yönlendirme Modeli" olarak formüle edilmiş; en kısa yol analizi ve benzetim teknikleriyle entegre edilerek sistematik bir çözüm geliştirilmiştir. Elde edilen sonuçlar, fiziksel yerleşimi değiştirmeden sadece hat yönlerinin eniyilenmesinin operasyonel performans ve maliyet verimliliğinde önemli iyileşmeler sağladığını kanıtlamaktadır. En dikkat çekici gelişme, ortalama araç sefer süresinin 37,55 dakikadan 33,63 dakikaya düşürülerek yaklaşık %10,45 oranında azaltılmasıdır. Ayrıca, Çift-Kapılı üretim hattında %11,28, Çok-Kapılı üretim hattında ise %4,74 oranında seyahat mesafesi tasarrufu sağlanmıştır.

Anahtar Sözcükler: Fabrika İçi Lojistik Eniyilemesi, Şerit Yönü Konfigürasyonu, Ağırlıklı Çizge Modeli, Seyahat Mesafesinin Minimizasyonu

Factory 6 Optimization of the Material Distribution Network

Abstract

This project addresses a critical inefficiency in the internal logistics system of Factory 6 by focusing on the optimization of AGV lane directions based on actual material flow patterns. By formulating the problem as a Weighted Arc Orientation Model (WAOM) and integrating it with shortest-path analysis and simulation, a systematic solution has been developed. The results demonstrate that optimizing lane directions alone, without altering the physical layout, can lead to significant improvements in both operational performance and cost efficiency. Most notably, the optimized configuration achieves a reduction of approximately 10.45% in average AGV round-trip time, decreasing from 37.55 minutes to 33.63 minutes, alongside travel distance reductions of 11.28% and 4.74% for the Double-Door and Multi-Door production lines, respectively.

Keywords: Intralogistics Optimization, Lane Direction Configuration, Weighted Arc Orientation Model, Travel Distance Minimization

6.1 Beko and Problem Identification

6.1.1 Company Description

Beko operates a specialized refrigerator manufacturing facility at Factory 6 in Eskişehir. The internal logistics network currently utilizes one AGV waiting station, three supermarkets, and two distinct production lines (Double-Door and Multi-Door) to manage material flow. Despite using a multi tiered trigger system and a hybrid model of manual and automated processes, the facility faces significant efficiency hurdles. High product variance on these assembly lines introduces supply complexities, making logistical precision essential for maintaining steady production.

6.1.2 Current System Analysis and Problem

The operational data provided by Beko describes the material handling activities within Factory 6, focusing specifically on AGV operations for the double-door and multi-door production lines. Currently, the factory utilizes a fleet of 13 active AGVs responsible for 40% of material movements, supplemented by an Automated Storage and Retrieval System (ASRS) and manual forklift operations. The material replenishment process is governed by a multi-level trigger system where a 90 minute stock alert moves parts from the ASRS to the supermarkets, and a 45 minute alert triggers AGV transport to specific workstations.

Detailed analysis reveals that the core problem is a suboptimal lane direction configuration that does not reflect historical flow patterns, leading to inefficient routing decisions. This misalignment forces vehicles to follow indirect routes to frequently visited stations, which results in the underutilization of the fleet and unnecessary total travel distance. These inefficiencies manifest as production bottlenecks where workstations wait for parts, as well as increased operational costs driven by higher energy consumption and accelerated mechanical wear on the AGVs. Furthermore, the existing road directions fail to provide a coordinated flow for the planned expansion to 23 AGVs and the updated layout featuring three distinct supermarkets.

6.2 Proposed Solution Strategy

The proposed strategy focuses on redesigning the lane-direction configuration of Factory 6's internal logistics network to minimize total AGV travel distance and enhance vehicle utilization. By formulating a Weighted Arc Orientation Model (WAOM), the project determines the optimal direction for each road segment based on the frequency of material cycles between the waiting area, three supermarkets, and production stations. This approach ensures that the lanes with the highest traffic volume are assigned the most direct and efficient directions, effectively tailoring the network to actual operational needs (Le-Anh and De Koster, 2006).

6.2.1 Critical Assumptions

The model is grounded in the updated factory layout provided in February, which incorporates a three-supermarket configuration. It assumes that all AGV cycles follow a fixed sequence: departing from the AGV waiting area (D), traveling to a designated supermarket for pickup, proceeding to a production station for delivery, travelling to crate handling area and returning to the waiting area (D) to receive new orders. Furthermore, demand weights for each station are derived from a detailed data analysis of the Bill of Materials (BOM) to accurately reflect the intensity of visits to each point.

6.2.2 Solution Approach

The solution approach is entirely data-driven, utilizing a mathematical optimization model implemented in Python with PuLP Solver to determine lane orientations. To capture demand intensity, a station weight parameter (w_r) is calculated based on station visit frequency data (Refer to Figure 6.1), representing exactly how often each location is visited by the AGVs. The specific demand frequencies used to feed the model, categorized by the Double-Door (DD), Multi-Door (MD), and Glass-Vegetable Drawer (CS) sections.

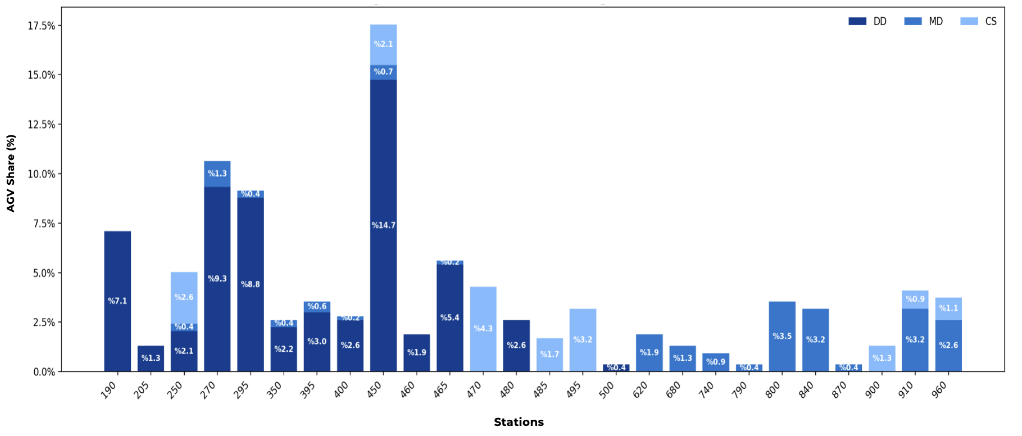

Figure 6.1: Frequency of each station from each Supermarket

The resulting optimized configuration is verified through shortest-path calculations using Dijkstra's algorithm and validated against historical performance to ensure the new directions provide a measurable improvement over the current layout.

6.2.3 Mathematical Model

Sets

- D: Set of AGV waiting area nodes.
- SM: Set of Supermarket pickup nodes.
- S: Set of production line stations.
- $N = D \cup SM \cup S$: Set of all nodes.
- $E \subseteq \{\{i,j\} \mid i,j \in N\}$: Set of undirected physical road segments.
- $A = \{(i,j) \mid \{i,j\} \in E\}$: Set of candidate directed arcs.
- $R \subseteq N \times N$: Set of required Origin-Destination pairs (a,b), where a is the origin and b is the destination. The set is defined as:

$$R = \{(d, sm) \mid d \in D, sm \in SM\} \cup \{(sm, s) \mid sm \in SM, s \in S\} \cup \{(s, d) \mid s \in S, d \in D\}$$

Parameters

- c_{ij}: Travel distance on directed arc $(i,j) \in A$.
- w_{ab}: Demand weight for required pair $(a,b) \in R$.
- M: A sufficiently large constant.

Decision Variables

$$X_{ij} = \begin{cases} 1, & \text{if the road is open in direction } i \to j, \\ 0, & \text{otherwise,} \end{cases} \qquad \forall (i,j) \in A$$

$$F_{ij}^{ab} = \begin{cases} 1, & \text{if the path from } a \text{ to } b \text{ uses arc } i \to j, \\ 0, & \text{otherwise,} \end{cases} \quad \forall (i,j) \in A,\ \forall (a,b) \in R$$

Objective Function

$$\min \sum_{(a,b)\in R} \sum_{(i,j)\in A} w_{ab}\, c_{ij}\, F_{ij}^{ab}$$

Constraints

$$X_{ij} + X_{ji} = 1 \quad \forall \{i,j\} \in E$$

$$\sum_{j:(i,j)\in A} F_{ij}^{ab} - \sum_{j:(j,i)\in A} F_{ji}^{ab} = \begin{cases} 1, & \text{if } i = a \\ -1, & \text{if } i = b \\ 0, & \text{otherwise} \end{cases} \qquad \forall i \in N, \forall (a,b) \in R$$

$$F_{ij}^{ab} \le M X_{ij} \quad \forall (i,j) \in A, \forall (a,b) \in R$$

$$X_{ij} \in \{0,1\} \quad \forall (i,j) \in A$$

$$F_{ij}^{ab} \in \{0,1\} \quad \forall (i,j) \in A, \forall (a,b) \in R$$

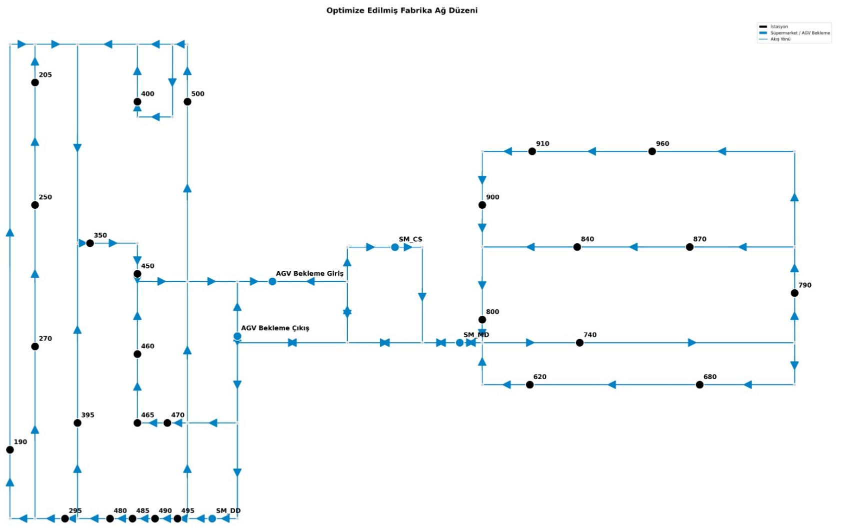

Figure 6.2: Output From the Mathematical Model

6.3 Verification

The verification phase was carried out to confirm the correctness and reliability of the proposed optimization model, as well as to ensure the proper implementation of the Arena simulation model employed in the subsequent validation stage under real factory conditions.

First, the mathematical optimization model and its Python implementation, solved using Gurobi, were tested using both simplified and expanded network structures. In reduced cases where the optimal solution was known, the model correctly identified the minimum-cost paths; in more complex scenarios, it maintained consistent behavior, demonstrating robustness. As part of the constraint feasibility test, the solution was systematically evaluated after solving the optimization model to ensure that all constraints were satisfied, including valid direction assignments, binary decision variables, and flow conservation conditions. After verifying the correctness of the optimization model, its output was used as an input to a Python-based implementation of Dijkstra's algorithm, which computes the shortest path distances between all node pairs. These distances were then used as inputs for the Arena simulation model. The Dijkstra algorithm was verified by comparing its outputs with manually calculated shortest paths, which matched exactly, confirming accuracy. Additional tests further showed that the algorithm correctly respects direction constraints and produces consistent results across repeated runs. The Arena simulation model, built to represent factory operations under both current and optimized lane directions, was verified through controlled experiments to ensure its correctness before use in the validation phase. A distance scaling test confirmed that travel times vary proportionally with distance, while sensitivity, extreme value, and replication tests demonstrated the model's reliability and robustness. Overall, the verification results confirm that all components of the model are correctly implemented and internally consistent.

6.4 Validation

To evaluate the accuracy and real-world applicability of the proposed solution, a validation process was conducted using both simulation and real system data. In the first phase, historical AGV data was divided into 30 batches and average round-trip times were calculated to form a representative dataset, which was then replicated in the Arena simulation model (See Figure 6.4) to obtain comparable simulated results; these two datasets were compared using a paired t-test at a 95% confidence level, and the results showed no statistically significant difference between historical and simulated data, confirming that the simulation model accurately represents real system behavior. In addition to statistical validation, performance comparisons between the current and optimized layouts demonstrated clear improvements, as both real data analysis and simulation results indicated reductions in total travel distance and average travel time under the optimized configuration.

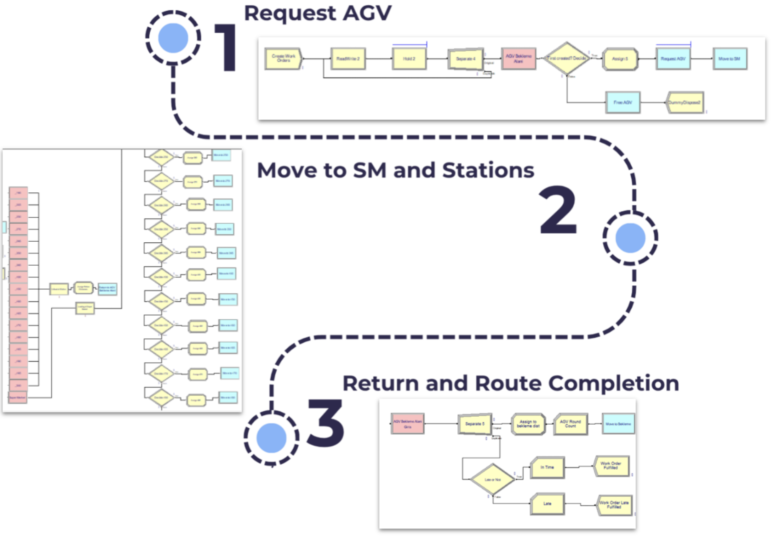

Figure 3: Arena Simulation Model

Table 6.1: Validation Results – Time Comparison

Scenario	**Average (minutes)**
Current Layout	37.55 minutes
Optimized Layout	33.63 minutes

Table 6.1 summarizes the impact of the proposed model on the average cycle time of the AGVs. By implementing the optimized lane direction configuration, the average travel time per cycle is reduced from 37.55 minutes to 33.63 minutes.

Table 6.2: Validation Results – Distance Comparison

Road Directions	**Total Distance (meters)**
with Current Road Directions- DD	19,046,080.00
with Optimized Road Directions- DD	17,119,360.00
with Current Road Directions- MD	536,948.00
with Optimized Road Directions- MD	511,504.00
Improvement- DD	11.28%
Improvement- MD	4.74%

Table 6.2 presents the total travel distance comparison between the current and optimized road direction configurations for both the Double-Door (DD) and Multi-Door (MD) production lines. These results confirm that the optimized lane direction configuration delivers measurable distance reductions across both production lines without any physical modifications to the factory layout.

In the second phase, field observations and company feedback were incorporated to assess practical applicability, and the company confirmed that the proposed lane directions and overall approach align with operational expectations and are suitable for implementation. Overall, the validation results demonstrate that the model is both accurate and practically applicable, providing meaningful operational improvements.

6.5 Decision Support System

To enhance the practical applicability of the proposed solution and ensure its usability in real operational settings, a Decision Support System (DSS) was developed to automate the optimization process and support future decision-making. The system is designed to work directly with company data, where users first upload an Excel file containing trip and station information together with the AutoCAD layout file of the factory (See Figure 6.3). Based on the uploaded data, the DSS automatically calculates station visit frequencies and constructs a frequency matrix, which is then used as demand weights in the mathematical optimization model. Using this input, the system runs the optimization model and generates the optimal lane direction configuration for the given layout. The final output is presented as an updated factory layout with optimized road directions (See Figure 6.2), allowing the company to directly visualize and evaluate the improved system. This structure enables the DSS to be flexible and reusable for different layouts and demand conditions, eliminating the need for manual model implementation. As confirmed by the company, the system is planned to be actively used in future operations, supporting more efficient intralogistics planning and data-driven decision-making.

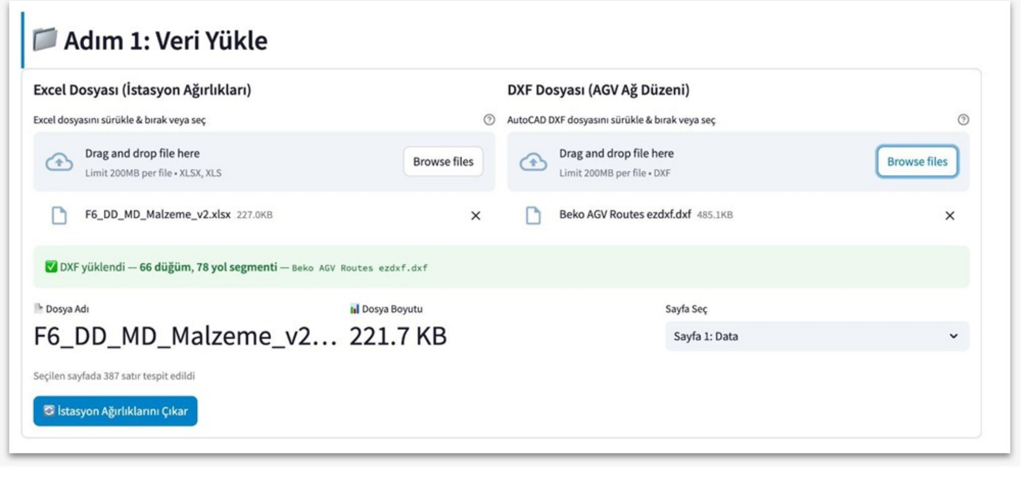

Figure 6.3: Decision Support System Interface

A "README" file was also provided to help the user navigate the decision support system and standardize the data in order to get the intended outcomes.

6.6 Benefits to the Company

The proposed solution provides significant operational and financial benefits to the company by improving the efficiency of the internal logistics system without requiring any physical changes to the factory layout. By optimizing lane directions based on actual material flow, the system reduces unnecessary AGV movements, leading to a measurable decrease in total travel distance and average transportation time. As demonstrated through both simulation and real data analysis, the optimized configuration results in a reduction of travel distance in both production lines and an overall decrease of approximately 10% in average travel time. These improvements directly contribute to increased system efficiency, reduced workstation idle times, and better utilization of the AGV fleet. In addition to operational gains, the company benefits from reduced energy consumption, resulting in annual energy savings of approximately 8,631.89 kWh and a corresponding cost reduction of around 36,685 TL, representing an improvement of about 11%. Furthermore, the integration of the Decision Support System enables the company to continuously adapt the model to changing layouts and demand conditions, supporting long-term, data-driven decision-making and eliminating the need for manual optimization efforts.

6.7 Conclusion

This project addresses a critical inefficiency in the internal logistics system of Factory 6 by focusing on the optimization of AGV lane directions based on actual material flow patterns. By formulating the problem as a Weighted Arc Orientation Model and integrating it with shortest-path analysis and simulation, a data-driven and systematic solution has been developed. The results demonstrate that optimizing lane directions alone, without altering the physical layout, can lead to significant improvements in both operational performance and cost efficiency. Through a rigorous verification and validation process, the model has been shown to be both internally consistent and representative of real system behavior. The observed reductions in travel distance are complemented by a substantial decrease in average cycle time; which dropped from 37.55 to 33.63 minutes, a 10.44% improve- ment that directly enhances fleet throughput. This reductions in travel distance, transportation time, and energy consumption highlight the prac- tical impact of the proposed approach. Furthermore, the development of a Decision Support System ensures that the solution is not limited to a single implementation but can be continuously adapted to future changes in layout and demand conditions. Overall, the project provides a scalable and sustainable framework for improving intralogistics operations, offering both immediate benefits and long-term strategic value to the company.

Bibliography

Le-Anh, T. and R. De Koster (2006, May). A review of design and control of automated guided vehicle systems. *European Journal of Operational Research 171*(1), 1–23.

7 İllikit Menkul Kıymetleri Getiri Eğrisi ile Değerleme Modeli

Emeklilik Gözetim Merkezi

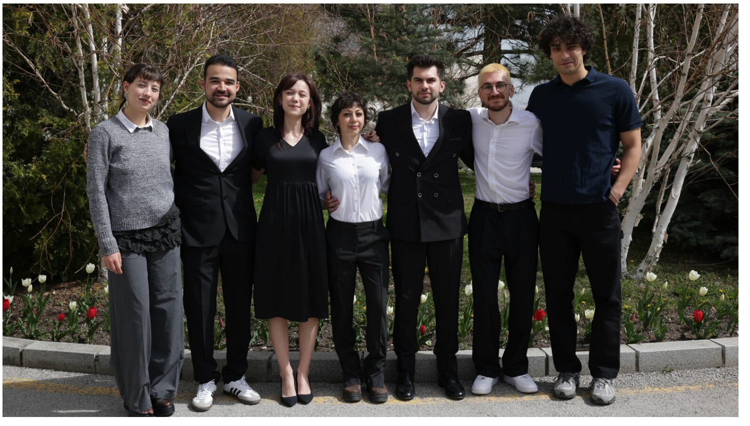

Proje Ekibi
Şevval Arslan, Ali Emre Bayraktar, İsmail Tuna Çakmak,
Tarık Emir Haliloğlu, İsmail Metlioğlu, Azra Yavaşcan, Irmak Yılmaz

Şirket Danışmanları
Hicran Aktaş Coşar, Barış Çakmak,
Yasemin Yüksel Çuhadar Kütle,
Ege Kaçar ve Samet Özben

Akademik Danışman
Doç. Dr. Çağın Ararat
Endüstri Mühendisliği Bölümü

Özet

Bu proje, Türkiye'deki özel emeklilik sisteminde likit olmayan sabit getirili menkul kıymetlerin değerlemesindeki tutarsızlıkları ele almaktadır. Mevcut yöntemlerin eski işlem fiyatlarına dayanması, piyasa koşullarını yansıtmayan sistematik fiyat sapmalarına yol açmakta; yapısal bir kilitlenme etkisi yaratmakta ve katılımcıların mevcut piyasa koşullarının sunduğu maksimum getiriyi elde etmesini engellemektedir. Bu soruna çözüm olarak, Nelson–Siegel–Svensson (NSS) modeli temelli bir getiri eğrisi çerçevesi önerilmektedir. Bu yaklaşım, tüm menkul kıymetlerin ortak bir yöntemle değerlenmesini sağlarken öznel uygulamaları, adaletsiz servet transferini ve arbitraj fırsatlarını azaltmaktadır.

Anahtar Sözcükler: Getiri Eğrisi, Sabit Getirili Menkul Kıymetler, Özel Emeklilik Sistemi, İllikidite, Likidite, Değerleme, Nelson Siegel Svensson Modeli, Spot Oranı

Yield Curve–Based Valuation Model for Illiquid Securities

Abstract

This project addresses valuation inconsistencies in illiquid fixed-income securities within Türkiye's private pension system. Existing methods rely on stale transaction prices, producing systematic mispricing that discourages portfolio rebalancing through a structural lock-in effect and prevents individual pension system participants from receiving the maximum return that current market conditions can offer. To address this, a yield curve–based framework is proposed, based on the Nelson–Siegel–Svensson (NSS) model, which provides a transparent and systematic representation of the term structure of interest rates. By valuing all securities under a common methodology, the framework reduces subjective inputs, eliminates arbitrage opportunities, and ensures fair and comparable outcomes across the pension system.

Keywords: Yield Curve, Fixed-Income Securities, Private Pension System, Illiquid, Valuation, Nelson-Siegel-Svensson Model, Liquid, Spot Rate.

7.1 Company Information

The Pension Monitoring Center (PMC) is the central regulatory body responsible for monitoring, data management, and coordination of Türkiye's private pension system. Operating under the authority of the Ministry of Treasury and Finance, PMC plays a key role in maintaining the stability, transparency, and reliability of the system (Emeklilik Gözetim Merkezi, 2025).

The private pension system, formally known as the Individual Pension System (IPS), was established in 2003 as a long-term savings and investment mechanism. Its primary purpose is to supplement public pension income by providing individuals with an additional source of retirement savings. PMC collects and maintains detailed records on participants, pension funds, contracts, and financial transactions, and delivers reporting support to relevant regulatory authorities, including the Insurance and Private Pension Regulation and Supervision Authority and the Capital Markets Board of Türkiye.

Through its Fund Monitoring Platform (FMP), PMC conducts daily oversight of pension funds and verifies the accuracy and consistency of asset valuations. The system currently serves approximately 17.5 million participants across 382 active pension investment funds, of which 347 contain fixed-income securities (FIS) such as government and corporate bonds. Given the scale of the system and the widespread presence of these instruments, accurate and consistent valuation of FIS is essential for participant fairness, transparent fund reporting, and effective regulatory oversight.

7.2 Current System and the Problem

The current valuation framework permits two different methods for pricing fixed-income securities (FIS). Although these methods differ in their calculation procedures, both rely on the most recent observed trade price to compute an internal rate of return (IRR), which is then used to discount future cash flows. As a result, both approaches are heavily dependent on the last recorded transaction price and may fail to capture current market conditions, particularly for securities that trade infrequently.

The current valuation framework relies on last trade prices, which creates significant distortions when market conditions change. Bonds purchased in low-interest-rate environments are priced at higher levels at lower market interest rates, but as the interest rates increase, the system continues to use these outdated prices as a reference. This results in systematic overvaluation, as the decline in the time value of money is not properly reflected. The problem is more pronounced for illiquid securities, where the absence of recent trades prevents price correction. Consequently, valuations deviate from market-consistent levels and reduce comparability across funds. This mispricing creates a structural lock-in effect: fund managers are discouraged from selling overvalued securities, as doing so would immediately reduce reported fund values. This behavior further suppresses trading activity, deepening market illiquidity over time. Ultimately, pension system participants bear the cost of this dynamic, earning systematically lower returns than what current market conditions would otherwise support.

In addition, differences in valuation timing, methodology, and discretionary inputs across fund management companies produce inconsistent pricing of identical securities. This lack of standardization reduces cross-fund comparability and creates conditions conducive to arbitrage and unfair wealth transfers between participants.

The central problem addressed in this project is therefore the absence of a standardized, market-consistent valuation framework for illiquid FIS. A model-based approach is needed to reduce dependence on stale transaction prices, restore market efficiency, and ensure that pension system participants receive fair returns that accurately reflect current market conditions.

7.3 Model and Proposed System

To address the valuation inconsistencies identified in the current system, a yield curve–based valuation framework is proposed. The proposed framework is developed in line with international best practices, supporting Türkiye's alignment with the International Financial Reporting Standards (IFRS), while remaining consistent with the valuation approaches adopted by institutions such as the European Central Bank (ECB). A yield curve represents the relationship between interest rates and maturities and provides a continuous term structure that can be used to discount future cash flows consistently across all securities.

Within this framework, the Nelson-Siegel-Svensson (NSS) model is selected to represent the term structure of interest rates. The NSS model provides a flexible yet parsimonious functional form that captures the level, slope, and curvature of the yield curve. Its widespread adoption by central banks and regulatory institutions, together with its ability to generate smooth and interpretable curves under limited data availability, makes it particularly suitable for markets. The base yield curve is constructed exclusively from liquid fixed-rate government bonds, ensuring that the term structure reflects reliable and current market information.

7.3.1 Mathematical Model

The mathematical framework consists of three main components: cash flow construction, yield curve modeling, and parameter estimation. For each bond, future cash flows are constructed based on its contractual characteristics, including maturity, coupon rate, and payment frequency (see Appendix 7.A.1). The time to each payment is calculated using a year fraction convention, ensuring consistency in maturity representation.

The term structure of interest rates is modeled using the NSS specification (see Appendix 7.A.2), which represents yields as a smooth function of maturity and captures the level, slope, and curvature of the yield curve. The model parameters are estimated by minimizing the sum of squared errors (SSE) between observed market prices and model-implied bond prices (see Appendix 7.A.5). To improve robustness, the optimization is performed over liquid government bonds only, defined as those with recent trading activity within a specified liquidity threshold. Maturity-sensitive weighting is also incorporated to prevent long-term bonds from disproportionately influencing the calibration. Once calibrated, the resulting yield curve is used to compute discount factors (see Appendix 7.A.3) and value all securities in the portfolio (see Appendix 7.A.4).

7.3.2 Proposed System

The proposed system introduces a standardized and automated valuation framework to replace the fragmented structure of current practice. By relying on a common yield curve as the reference, the system ensures that securities with identical cash flow structures are valued consistently across all fund management companies. The automated pipeline integrates data input, parameter estimation, and valuation within a single structure, reducing operational errors while improving auditability and reproducibility. The modular design of the framework allows the system to be extended to different instrument types without altering the core structure, maintaining both consistency and flexibility. Overall, the proposed system provides a scalable, transparent, and regulator-aligned solution for the valuation of illiquid FIS.

7.4 Validation of the Approach

To evaluate the performance of the proposed framework, a controlled validation study is conducted using government bonds with known market prices. These bonds are deliberately treated as illiquid by withholding their most recent prices from the model, and instead supplying older transaction prices as inputs to the existing valuation methods. This setup allows for a direct and observable comparison between the two approaches under identical conditions, as the true market price serves as the reference for error measurement.

Valuation accuracy is measured using Mean Absolute Percentage Error (MAPE), which expresses the average deviation between model-implied and observed market prices as a percentage. MAPE is selected as the primary metric because it is scale-independent and allows consistent comparison across bonds with different price levels.

The results indicate that the proposed NSS-based framework produces substantially lower valuation errors compared to the existing methods. This improvement stems from the model's reliance on a yield curve calibrated from current market data, rather than on outdated last-trade prices. Detailed numerical results and bond-level comparisons are in Section 7.6.

7.5 Integration and Implementation

The proposed framework is implemented as a web-based decision support tool designed to standardize the valuation of FIS within pension fund management as seen in Figure 7.1. The system is structured to be integrated into the Fund Monitoring Platform (FMP) while also supporting standalone use by individual fund management companies.

The implementation process begins with data input. Users upload a bond cashflow file in a predefined Excel format containing bond-level attributes such as coupon structure, maturity, and trading history. Following the data upload, users specify the valuation date and liquidity threshold.

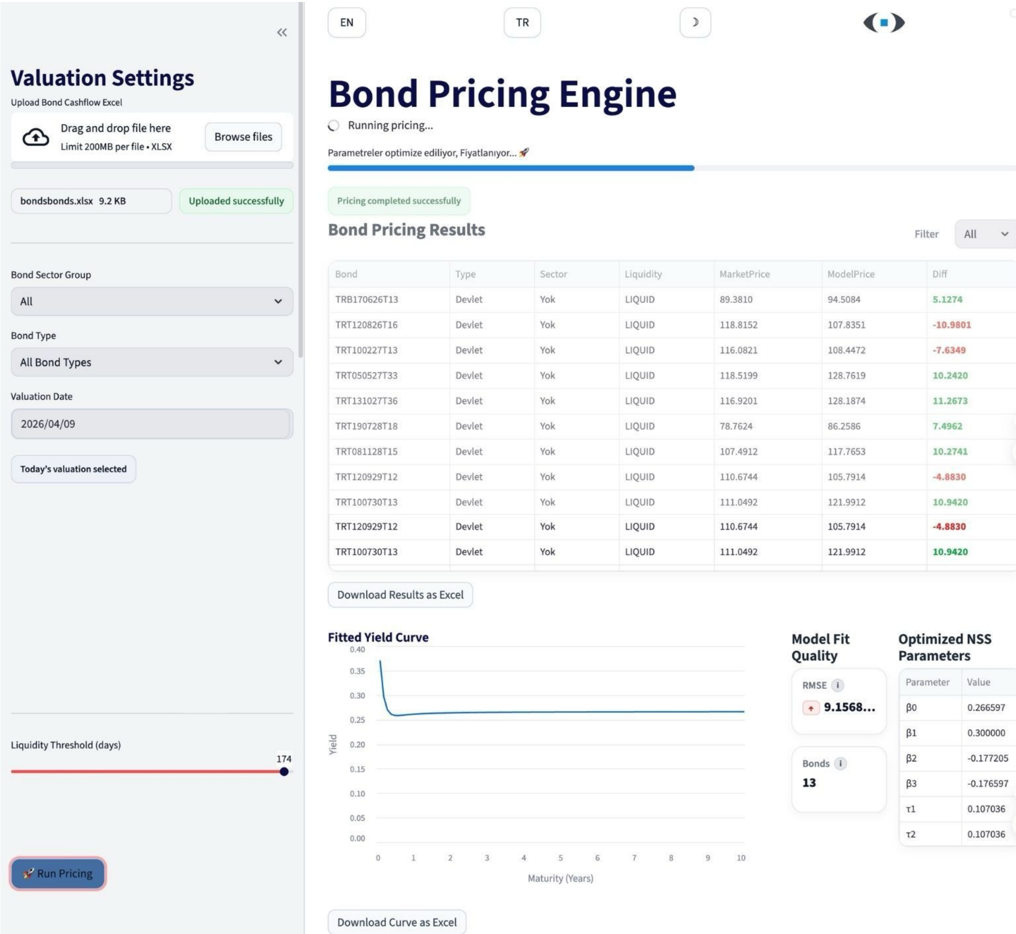

Figure 7.1: Decision Support System

The liquidity threshold determines which bonds are considered sufficiently active for yield curve calibration.

Once the inputs are confirmed, the system automatically calibrates the NSS model using the selected liquid government bonds and constructs the base yield curve. The calibrated curve is then applied to value all securities in the dataset. As outputs, the system provides the optimized NSS parameters, a visualization of the fitted yield curve, and the calculated bond prices, all of which can be exported in Excel format for further review and reporting purposes.

This design ensures reproducibility, scalability, and auditability, making the system well-suited for regulatory oversight and systematic implementation across the pension fund ecosystem.

7.6 Benchmarking

The proposed NSS-based valuation framework is benchmarked against the current system using 35 bonds traded in 2026. The comparison is based on valuation accuracy, measured by Mean Absolute Percentage Error (MAPE). Results show a clear improvement in performance. The average MAPE of the current system is 9.61%, whereas the NSS model reduces this to 6.98%, corresponding to an improvement of approximately 27%.

In addition, forward comparisons using updated market prices indicate that NSS-based valuations are closer to subsequent realized prices. This indicates that the model not only improves valuation accuracy but also captures the underlying term-structure dynamics more effectively. Overall, the results demonstrate that the NSS framework provides more accurate, stable, and economically consistent valuations compared to the existing approach, especially for illiquid fixed-income securities.

7.7 Benefits to the Company

The proposed framework delivers measurable benefits to both PMC and the broader pension ecosystem. The NSS-based approach reduces average valuation error by approximately 27% compared to existing methods, providing a quantifiable improvement in pricing accuracy.

By ensuring that securities are priced at market-consistent values, fund managers are no longer discouraged from rebalancing their portfolios in response to changing market conditions. This restores trading activity, improves overall market liquidity, and most importantly, ensures that IPS participants receive the maximum return that current market conditions can offer.

From PMC's perspective, the framework reduces reliance on fund-reported inputs and manual valuation practices, lowering the risk of operational errors and misreporting. Its transparent and reproducible structure strengthens auditability and supports more effective regulatory oversight, while eliminating cross-fund pricing discrepancies and preventing unfair wealth transfers between participants.

7.8 Conclusions and Future Work

This project proposes a yield curve–based valuation framework to address the structural limitations of current valuation practices in Türkiye's private pension system. The existing reliance on stale transaction prices produces systematic mispricing of illiquid fixed-income securities, creating a lock-in effect that suppresses market activity and prevents IPS participants from receiving the maximum return that current market conditions can offer.

The proposed NSS-based framework replaces this approach with a market-consistent term structure, achieving a 27% reduction in average valuation error while establishing a standardized, transparent, and reproducible valuation process across the pension system.

Future work may focus on automating the calibration of the liquidity threshold using data-driven approaches, adding a sector-specific risk premium for corporate bonds on top of the base curve to account for credit-related characteristics, and developing a more structured methodology for estimating sector-specific credit and liquidity spreads for corporate bond valuations.

Bibliography

Emeklilik Gözetim Merkezi (2025). Emeklilik gözetim merkezi – bireysel emeklilik sistemi. https://www.egm.org.tr/. Accessed: 2025-11-01.

Appendices

7.A Mathematical Model

7.A.1 Time and Cash Flow Structure

- t_0: valuation date
- $t_{i,j}$: j-th payment date of bond i
- $\tau_{i,j} = \text{YearFraction}(t_0, t_{i,j})$
- n_i: number of cash flows of bond i

$$CF_{i,j} = \begin{cases} \frac{c_i}{f_i} F, & j = 1, \dots, n_i - 1, \\ \frac{c_i}{f_i} F + F, & j = n_i \end{cases}$$

7.A.2 NSS Yield Curve Model

$$y_\theta(\tau) = \beta_0 + \beta_1 \frac{1 - e^{-\lambda_1 \tau}}{\lambda_1 \tau} + \beta_2 \left(\frac{1 - e^{-\lambda_1 \tau}}{\lambda_1 \tau} - e^{-\lambda_1 \tau} \right) + \beta_3 \left(\frac{1 - e^{-\lambda_2 \tau}}{\lambda_2 \tau} - e^{-\lambda_2 \tau} \right)$$

- β_0: Level
- β_1: Slope

- β_2: Curvature (medium)
- β_3: Curvature (long)
- λ_1: Decay (short–medium)
- λ_2: Decay (long)
- θ: Vector of model parameters

7.A.3 Discounting

$$d_{i,j}(\theta) = e_{\theta}^{-y(\tau_{i,j})\tau_{i,j}}$$

7.A.4 Bond Pricing

$$P_i^{\mathrm{model}}(\theta) = \sum_{j=1}^{n_i} CF_{i,j}\, d_{i,j}(\theta)$$

7.A.5 Calibration

$$\min_{\theta} \sum_{i \in \mathcal{L}} w_i \left(P_i^{\mathrm{model}}(\theta) - P_i^{\mathrm{market}} \right)^2$$

- $\theta = (\beta_0, \beta_1, \beta_2, \beta_3, \lambda_1, \lambda_2)$
- $\mathcal{L}$: liquid bonds
- w_i: weights
- P_i^{market}: observed price

İthal Ham Madde Tedarik Süreçlerinin Eniyileştirilmesi | 8

Eti Gıda

Proje Ekibi
Erkin Aşan, Arda Bayrakçı, Azra Ceren Durmuş, Mert Ertunga Ekizce, Ahmet Efekan Ersoy, Ömer Kıvanç Tanrıverdi, Sahra Yalçın

Şirket Danışmanı
Irmak Çelikkol Çalışkan
Üretim Planlama Yöneticisi

Akademik Danışman
Doç. Dr. Ayşe Selin Kocaman
Endüstri Mühendisliği Bölümü

Özet

Bu proje, teslim süresi belirsizliği, değişken talep ve operasyonel karmaşıklığın neden olduğu yüksek stok maliyetlerini, acil sevkiyatları ve planlama zorluklarını azaltmayı amaçlamaktadır. İthal ham madde tedarik süreçlerinin daha etkin ve sistematik şekilde yönetilmesi hedeflenmektedir. Bu amaçla, belirsizlik altında karar almayı destekleyen veri odaklı bir tedarik planlama yaklaşımı geliştirilmiştir. Önerilen model, çok dönemli bir yapı içerisinde Örnek Ortalama Yaklaşımı ve Kayan Ufuk Sezgiseli kullanarak esnek sipariş kararları üretmektedir. Doğrulama çalışmaları, maliyetlerin ağırlıklı ortalamasında %13,34 iyileşme ve stokların ağırlıklı ortalamasında %14,80 azalma elde edildiğini göstermektedir. Sistem, günlük planlama süreçlerinde daha hızlı, güvenilir ve tutarlı kararlar alınmasına katkı sağlamaktadır.

Anahtar Sözcükler: Tedarik Süresi Belirsizliği, Envanter ve Tedarik Planlaması, Örnek Ortalama Yaklaşımı, Kayan Ufuk Sezgiseli, Dinamik Programlama, Tedarik Zinciri Eniyilemesi.

Optimization of Imported Raw Material Procurement Processes

Abstract

This project aims to improve the management of imported raw material procurement processes by reducing high inventory costs, emergency shipments, and planning difficulties caused by lead-time uncertainty, demand variability, and operational complexity. To address these challenges, a data-driven procurement planning approach is developed to support decision-making under uncertainty. The proposed model generates adaptive ordering decisions within a multi-period structure using a Sample Average Approximation Rolling Horizon (SAA-RH) heuristic. Validation results show that an overall weighted cost improvement of 13.34% and a 14.80% weighted average inventory reduction. The system enables faster, more reliable, and consistent decision-making in daily procurement operations.

Keywords: Lead-Time Uncertainty, Inventory and Procurement Planning, Sample Average Approximation Rolling Horizon Heuristic, Dynamic Programming, Supply Chain Optimization.

8.1 Company Description

Eti is one of the leading food manufacturing companies in Türkiye. Founded in 1962, the company produces a wide range of products, including biscuits, wafers, cakes, chocolates, cereals, baby food, and healthy snack alternatives. Eti operates several production facilities located in Eskişehir, Bozüyük, and Konya, and exports to over 100 countries. Due to the scale and complexity of its production operations, ensuring the timely and reliable supply of raw materials is critical. In particular, the procurement of imported raw materials such as cocoa, oils, and food additives, sourced from international suppliers via multiple transportation modes, plays a significant role in supporting production continuity and operational efficiency.

8.2 System Analysis

The imported raw material procurement system at Eti has been managed by the Production Planning Department since September 2024. Despite this structural change, the system still relies heavily on manual intervention, limiting overall efficiency. Although the process begins with MRP execution, final order quantities are determined uingExcel-based data extracted from SAP, resulting in a largely manual process.

Procurement decisions are limited to predefined supplier agreements and predefined transportation modes (seaway, roadway), along with minimum order quantities (MOQ). While seaway is usually chosen due to its lower cost, the lack of a structured planning approach often leads to reactive decisions, where expensive roadway transportation is used only when problems arise, rather than being planned in advance.

Lead times constitute a major source of uncertainty. In practice, planners rely on conservative buffering strategies to reduce stockout risk, which directly affects inventory levels and ordering decisions. As a result, the system operates in a predominantly deterministic manner, without pipeline visibility or explicit modeling of uncertainty. This limits its ability to respond effectively to variability and leads to suboptimal procurement decisions.

8.3 Problem Definition

Eti's imported raw material planning system is challenged by lead-time uncertainty, multi-supplier sourcing, and largely manual planning practices. The company manages 41 imported materials with varying and unpredictable delivery times, making it difficult to ensure consistent and reliable replenishment decisions (Kaplan, 1970). This variability directly affects the system's ability to maintain a stable balance between supply and demand.

To prevent production disruptions, the current system relies on conservative buffering strategies based on maximum observed lead times. While this approach reduces stockout risk, it leads to higher inventory levels and associated costs, as well as an increased risk of waste due to perishability. At the same time, unexpected delays or demand fluctuations force the system into costly emergency decisions, particularly the use of high-cost transportation options.

An ABC analysis reveals that a small number of high-value items account for approximately 80% of total annual consumption value, making planning decisions for these materials especially critical (Jacobs et al., 2011). As a result, inventory levels fluctuate between excess and shortage, creating misalignment between material availability and production needs. Overall, the core problem lies in the inability of the current system to manage uncertainty in a structured and balanced way, leading to inefficient cost and service outcomes.

8.4 Proposed Solution

8.4.1 Critical assumptions and data analysis

The model is developed under a set of simplifying assumptions to ensure tractability. It focuses on A-class raw materials identified through ABC classification, as these items account for the majority of total consumption value. Demand is treated as deterministic over the planning horizon, given the high accuracy of existing forecasts. The primary source of uncertainty is lead-time variability, which necessitates a stochastic solution approach. This is modeled using a scenario-based framework where historical observations are directly used as discrete scenarios instead of fitting a predefined distribution. Since materials are sourced from multiple suppliers, lead-time variability is estimated through a pooled variance approach, weighting supplier-specific variances by sample sizes. In addition, cross-orders, where a later order arrives earlier than a previous one, are excluded with company approval, as they are rare and have a negligible impact. Finally, transportation mode availability is assumed to be predetermined for each material based on supplier agreements, restricting decisions to feasible supplier–mode combinations.

8.4.2 Major constraints

The model incorporates operational constraints to ensure feasibility under lead-time uncertainty. Orders are subject to supplier-specific minimum order quantities (MOQ), requiring placement in integer multiples of predefined batch sizes. Binary decision variables restrict ordering to feasible supplier–transportation mode combinations. Inventory balance constraints ensure that demand is satisfied using available inventory whenever possible, while unmet demand is carried forward as backlog. Non-negativity constraints are imposed on inventory and backlog variables, and integrality and binary requirements are enforced on order quantities and decision variables.

8.4.3 Objectives

The objective of the following model is to minimize the total accumulated operational cost over the planning horizon. The objective function includes purchasing and transportation costs, setup costs, inventory holding costs, and backorder penalties. An additional penalty cost is defined for the use of emergency road transportation. This penalty is introduced to discourage the routine use of this mode and to ensure that it is applied only as a corrective measure when necessary. As a result, critical stock levels are protected without relying excessively on high-cost transportation alternatives.

8.4.4 Conceptual model

To address challenges posed by lead-time uncertainty and the complex structure of the procurement system, a sequential decision-making model is developed. The model relies on a state-dependent policy, in which decisions are made based on the current system state. The state representation includes the on-hand inventory level, unmet demand (backorders), and pipeline inventory that tracks all outstanding orders along with their remaining lead times. This structure ensures that both current system conditions and the delayed effects of past decisions are fully captured within the decision process. Uncertainty in the system is modeled through scenario-based representations of lead times. Instead of assuming a specific probability distribution, historical lead-time observations are directly used to construct a discrete set of possible scenarios (Huang, Kai and Ahmed, Shabbir, 2009). At each decision period, these scenarios represent different realizations of delivery times across transportation modes, allowing the model to capture real-world variability in supply processes. The combination of scenarios reflects the joint evolution of lead times and enables the evaluation of decisions under multiple possible future outcomes.

8.4.5 Mathematical model

The mathematical model minimizes expected total operational cost under lead-time uncertainty, with the full formulation provided in Appendix: Mathematical Model. Inventory dynamics are represented through pipeline-based state transitions, where outstanding orders evolve under stochastic lead times and arrivals update on-hand inventory (Disney et al., 2016).

The problem is formulated as a finite-horizon dynamic programming model, where the system state consists of inventory, backlog, and pipeline variables, and decisions are made sequentially over time. Safety stock constraints regulate the use of emergency transportation (Chopra and Meindl, 2016).

Computational experiments show that the runtime of the DP algorithm is mainly driven by the planning horizon, number of scenarios, and lead-time magnitude. Increasing the number of scenarios from 8 to 16 raises solution time from 44 seconds to over 1,600 seconds (nearly 38-fold increase), while higher lead times expand the pipeline state space, increasing computation time by up to 330%.

Therefore, runtime grows non-linearly with key problem dimensions, making the exact DP formulation computationally intractable for a 90-day horizon. Furthermore, a Sample Average Approximation Rolling Horizon(SAA-RH) heuristic approach is adopted, in which the problem is repeatedly solved over a limited look-ahead window. This provides a tractable approximation of the original DP model while allowing the system to adapt dynamically to new information.

8.5 Solution Method

8.5.1 Sample average approximation rolling horizon (SAA-RH) heuristic solution approach

The proposed solution method addresses the computational challenges of the exact dynamic programming model by employing an SAA-RH heuristic. The 90-day planning horizon is divided into daily decision periods, within which procurement decisions are updated sequentially based on the current system state, including on-hand inventory, backlog, and pipeline orders, as represented in Figure 8.1 (Manuel Díaz-Madroñero and Josefa Mula and Mariano Jiménez and David Peidro, 2017).

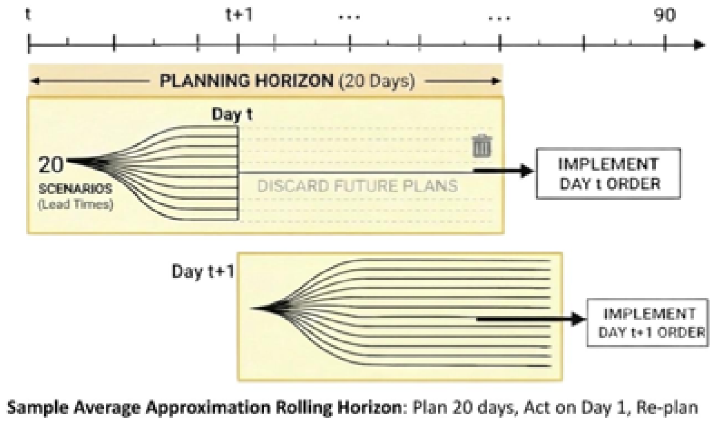

Figure 8.1: SAA-RH Solution Illustration

Lead-time uncertainty is captured through a finite set of scenario-based realizations derived from historical data for both road and sea transportation modes. Within the Sample Average Approximation framework, the model is solved over the scenario set, and decisions are determined by minimizing the sample average of total cost across scenarios. This ensures that the resulting policy is robust to lead-time uncertainty. In line with the SAA-RH approach, only the first-period decisions are implemented. The system state is then updated and carried forward to the next period, and the process is repeated throughout the planning horizon. This rolling structure produces a state-dependent and adaptive ordering policy while maintaining computational feasibility for daily use.

8.5.2 Implementation and software architecture

The solution framework is implemented in Python, integrating data preprocessing and the SAA-RH heuristic approach. This environment provides sufficient flexibility to handle large datasets and iterative solution procedures across multiple materials, while supporting repeated re-optimization at each decision period.

The mathematical model is formulated using Pyomo, which enables a structured representation of the objective function and operational constraints in an open-solver compatible form. The resulting optimization problems are solved using the HiGHS solver, providing an effective balance between computational performance and accessibility. To ensure practical runtimes, a 5% optimality gap is applied, allowing the model to generate reliable solutions within acceptable computational time for daily use.

8.6 Validation

The proposed model was validated using real operational data from Eti's supply chain through a pilot study focusing on high-impact raw materials, representing 67.4% of total inventory cost. The results show consistent cost improvements across all materials, as presented in Table 8.1, with reductions ranging between 10.41% and 41.89%, leading to an overall weighted cost improvement of 13.34%.

Table 8.1: Comparison of Cost, Service Level, and Inventory Improvements

Material ID	**Material Proportion in Total Cost (%)**	**Service Level (%)**	**Cost Improvement (%)**	**Inventory Level Improvement (%)**
RM09	41.5%	95.6% → 96.7%	10.41%	10.34%
RM04	13.72%	100.0% → 100.0%	34.97%	34.97%
RM06	6.78%	100.0% → 92.0%	41.89%	46.23%
RM05	5.4%	100.0% → 90.6%	25.57%	48.01%
TOTAL	67.4%		13.34% (Weighted)	14.80% (Weighted)

The pilot study also reveals a strategic trade-off in service levels. RM09 improved to 96.7%, while RM06 and RM05 were adjusted to 92.0% and 90.6%, respectively, to reduce overstocking. This indicates that the model reduces unnecessary inventory accumulation while maintaining service performance without significant deterioration.

Furthermore, the model achieved inventory reductions of 10.34% (RM09), 34.97% (RM04), 48.01% (RM05), and 46.23% (RM06). When weighted by material share, these improvements correspond to a 14.80% reduction in overall inventory levels. Overall, the results demonstrate effective inventory rationalization, leading to improved cost efficiency and more balanced procurement decisions. These findings confirm the practical applicability and robustness of the proposed decision support system.

8.7 Benchmarking and Benefits

A benchmarking analysis was conducted comparing the exact Dynamic Programming (DP) approach with the SAA-RH approach in terms of cost performance, computational efficiency, and operational applicability. The DP model provides a theoretical global optimum by anticipating future uncertainties and preventing backorders through proactive use of faster transportation. The SAA-RH approach resulted in solutions with an optimality gap between 7.45% and 12.77%, mainly due to its limited look-ahead structure. The detailed benchmarking results are presented in Table 8.2.

From a computational perspective, the DP approach requires approximately 94 minutes to solve a 30-day planning problem, whereas the RH system generates decisions in under a second (around 0.31 seconds). Sensitivity analysis further showed that extending the planning horizon and increasing demand scenarios led to cost improvements of up to 13.09% while maintaining negligible computation time. Beyond numerical performance, the system supports managerial decision-making by replacing manual judgment with data-driven purchasing policies and enabling structured supplier comparisons.

Table 8.2: Comparison of DP and Rolling Horizon Approaches

Horizon	**DP (Optimal) Expected Cost [TL]**	**RH (Simulation) Mean Cost [TL]**	**Optimality Gap [%]**	**DP Solution Time**	**RH Solution Time (Per Run)**
10	1,118,162.00	1,201,560.36	7.45%	700 sec	196.86 sec
20	2,101,875.00	2,370,295.83	12.77%	2840 sec	795 sec
30	3,096,542.00	3,479,868.42	12.38%	5656 sec	1583 sec

The results were obtained in a stochastic environment using the SAA-RH approach with 6 distinct lead-time scenarios ($N = 6$). The simulation incorporates multiple suppliers and transportation modes with varying lead-time distributions and cost profiles. Operating on a 5-day Planning Horizon, the system dynamically trades off between sourcing options to anticipate demand fluctuations. By optimizing across these scenarios and suppliers simultaneously, the model minimizes total operational costs while ensuring supply chain resilience.

Beyond numerical performance, the proposed system provides substantial operational and financial benefits. Reducing excess inventory improves cash flow and working capital utilization while also lowering the risk of production disruptions through better anticipation of delays. In addition, the model enhances decision transparency by quantifying cost–risk trade-offs and standardizing procurement decisions, reducing dependence on manual judgment. Unlike the current system, it incorporates full visibility of pipeline inventory and dynamically balances holding, shortage, and transportation costs under uncertainty. This framework automates the whole procurement process. Finally, its flexible structure allows easy adaptation across different operational settings, enabling scalable and data-driven procurement strategies.

8.8 Implementation and Pilot Study

The decision support system was implemented as an automated tool and integrated into Eti's procurement planning framework, and was delivered together with its installation files on March 16, 2026. The model operates with short computational times, enabling frequent updates without disrupting existing workflows. Users interact with the system through Excel-based outputs and a Streamlit-based interface shown in Figure 8.2, where recommended order quantities and projected inventory levels are presented in a clear and actionable format. Input data, including inventory levels, lead times, and pipeline information, is managed through standardized Excel templates.

Figure 8.2: DSS Home Page

Decisions are generated using scenario-based simulations to capture demand and supply uncertainty. The system supports both full-scale planning runs (covering all materials) and item-level analysis (for selected materials), allowing flexibility depending on planning needs. The outputs are structured as a chronological execution plan, enabling direct use in operational decision-making.

Following delivery, a pilot study was conducted. By April 6, 2026, the system had been fully applied to five materials in class A, and the results showed strong alignment with the company's actual historical decisions and timelines. This close correspondence increased confidence in the model and demonstrated its practical applicability. Feedback from training sessions and pilot usage indicated that the tool is intuitive, aligns well with existing planning practices, and has been well received by company stakeholders.

The development process was carried out in close collaboration with company experts through regular meetings and technical reviews, improving both usability and practical relevance. Preliminary results show that the system effectively supports data-driven procurement decisions by replacing manual approaches with a structured methodology. The model produces stable and feasible order plans under varying demand scenarios while maintaining service levels and controlling inventory risks. The weighted cost improvement analysis indicates a total potential impact of 13.34%.

8.9 Conclusion

The project met expectations at Eti and demonstrated strong potential for broader implementation by improving production planning, reducing costs, and enhancing decision consistency. A decision support system based on stochastic optimization and the Sample Average Approximation Rolling Horizon heuristic improved inventory control, reduced excess stock levels, and enabled more reliable and data-driven production planning decisions under uncertainty. The system achieved a weighted cost improvement of 13.34% and reduced inventory levels by 14.80%, across high-impact materials representing 67.4% of the total inventory cost. This was achieved while maintaining service performance and mitigating disruption risks by leveraging scenario-based modeling and stochastic optimization.

In addition, the system replaced manual judgment with a structured quantitative framework, increasing transparency in cost–risk trade-offs and standardizing decision processes. It also reduced managerial workload by automating planning, calculation, and reporting tasks, enabling faster and more efficient daily operations.

The system is easily adaptable across different materials and operational settings with minimal customization, providing a scalable and practical solution for procurement and inventory planning under uncertainty, managing cost–service trade-offs under uncertainty.

Bibliography

Chopra, S. and P. Meindl (2016). *Supply Chain Management: Strategy, Planning, and Operation* (6th ed.). Pearson.

Disney, S. M., A. Maltz, X. Wang, and R. D. Warburton (2016). Inventory management for stochastic lead times with order crossovers. *European Journal of Operational Research 248*(2), 473–486.

Huang, Kai and Ahmed, Shabbir (2009). The value of multistage stochastic programming in capacity planning under uncertainty. *Operations Research 57*(4), 893–904.

Jacobs, F., R. Chase, and R. Lummus (2011). *Operations and Supply Chain Management*. McGraw-Hill/Irwin series operations and decision sciences. McGraw-Hill Irwin.

Kaplan, R. S. (1970). A dynamic inventory model with stochastic lead times. *Management Science 16*(7), 491–507.

Manuel Díaz-Madroñero and Josefa Mula and Mariano Jiménez and David Peidro (2017). A rolling horizon approach for material requirement planning under fuzzy lead times. *International Journal of Production Research 55*(8), 2197–2211.

Appendix: Mathematical Model

Sets and indices

Symbol	**Explanation**
T	Set of planning periods, $T = \{1, \ldots, \|T\|\}$
I	Set of suppliers, $i \in I$
J	Set of transportation modes, $j \in J$
L	Set of historical lead times, $L = \{1, \ldots, \bar{L}\}$
S_t	System state at period t

Decision variables

Symbol	**Explanation**
$O_{i,j,t}$	Order quantity decided for supplier i via mode j in period t
$x_{i,j,t}$	Binary ordering decision; equals 1 if an order is placed, and 0 otherwise
$Y_{i,j,t}$	Number of MOQ batches ordered for pair (i, j) in period t
z_t	Binary activation of emergency transportation in period t
u_t	Decision vector at period t

Parameters and state variables

Symbol	Explanation
D_t	Deterministic demand in period t
$c_{i,j}$, $K_{i,j}$	Unit purchasing and fixed ordering costs for pair (i, j)
h, p, ρ	Unit holding, backorder, and emergency activation costs
$MOQ_{i,j}$, $\bar{Y}_{i,j}$	Minimum order quantity and maximum allowable multiple
$a_{i,j}$	Binary availability of mode j for supplier i
SS, M	Safety stock threshold and a sufficiently large constant
$L_{i,j,t}$	Realized stochastic delivery time for order placed in period t
I_t, B_t	Inventory level and backorder quantity at the end of period t
$U_{i,j,l,t}$	Quantity of outstanding orders with l periods remaining until delivery
$\tilde{U}_{i,j,l,t}$	Shifted pipeline inventory before new orders are assigned in period t
R_t	Total arrival quantity received at the beginning of period t
$V_t(S_t)$	Optimal value, or cost-to-go, function at period t
$C_t(S_t, u_t)$	Immediate cost incurred in period t at state S_t and decision u_t

Objective function

The objective is to minimize the expected total operational cost:

$$\min \ \mathbb{E}\left[\sum_{t\in T}\left(\sum_{i\in I}\sum_{j\in J}(c_{i,j}O_{i,j,t} + K_{i,j}x_{i,j,t}) + hI_t + pB_t + \rho z_t\right)\right]$$

Dynamic programming representation

The system state at period t is defined as

$$S_t = (I_{t-1}, B_{t-1}, U_{i,j,l,t-1} \ \ \forall i \in I,\ j \in J,\ l \in L)$$

The sequential decision-making process is solved within a dynamic programming framework:

$$V_t(S_t) = \min_{u_t}\left\{C_t(S_t, u_t) + \mathbb{E}[V_{t+1}(S_{t+1}) \mid S_t, u_t]\right\}, \quad \forall t \in T$$

Constraints

$$\tilde{U}_{i,j,l,t} = \begin{cases} U_{i,j,l+1,t-1}, & l = 1, \dots, \bar{L} - 1 \\ 0, & l = \bar{L} \end{cases} \qquad \forall i, j, l, t \tag{8.1}$$

$$U_{i,j,l,t} = \tilde{U}_{i,j,l,t} + \mathbf{1}\{L_{i,j,t} = l\}\, O_{i,j,t} \qquad \forall i, j, l, t \tag{8.2}$$

$$R_t = \sum_{i \in I} \sum_{j \in J} U_{i,j,1,t-1} \qquad \forall t \in T \tag{8.3}$$

$$I_t - B_t = I_{t-1} + R_t - D_t - B_{t-1} \qquad \forall t \in T \tag{8.4}$$

$$O_{i,j,t} = MOQ_{i,j}\, Y_{i,j,t} \qquad \forall i, j, t \tag{8.5}$$

$$x_{i,j,t} \leq Y_{i,j,t} \leq \bar{Y}_{i,j}\, x_{i,j,t} \qquad \forall i, j, t \tag{8.6}$$

$$x_{i,j,t} \leq a_{i,j} \qquad \forall i, j, t \tag{8.7}$$

$$I_{t-1} \leq SS + M(1 - z_t) \qquad \forall t \in T \tag{8.8}$$

$$x_{i,j,t} \leq M z_t \qquad \forall j \in J_{\text{road}},\ i \in I,\ t \tag{8.9}$$

$$O_{i,j,t},\ I_t,\ B_t,\ R_t \geq 0 \qquad \forall i, j, t \tag{8.10}$$

$$x_{i,j,t},\ z_t \in \{0, 1\} \qquad \forall i, j, t \tag{8.11}$$

$$Y_{i,j,t} \in \mathbb{Z}^+ \qquad \forall i, j, t \tag{8.12}$$

Constraint descriptions

- (8.1)–(8.2) define the pipeline state transition under stochastic lead-time realizations.
- (8.3) computes arrivals from the outstanding order pipeline.
- (8.4) enforces inventory and backlog flow balance across periods.
- (8.5)–(8.6) impose minimum order quantities and ordering activation logic.
- (8.7) ensures that only feasible supplier–mode combinationsis selected.
- (8.8)–(8.9) control the conditional activation of emergency orders under safety stock violations.
- (8.10)–(8.12) define variable domains and integrality restrictions.

9 | Kantar Operasyonlarının Optimizasyonu

Limak Çimento

Proje Ekibi
Deniz Bodur, Ahmet Melih Bülbül,
Mert Can İbri, Begüm Şimşek

Şirket Danışmanı
Ayça Küçükoğlu
Lojistik Departmanı Yöneticisi

Akademik Danışman
Prof. Dr. Ülkü Gürler
Endüstri Mühendisliği Bölümü

Özet

Limak Çimento fabrikalarında gerçekleştirilen bu çalışmada, yüksek araç trafiği nedeniyle kantar operasyonlarında oluşan verimsizlik problemi ele alınmıştır. Mevcut sistemde manuel veri girişi, sözlü iletişim ve fiziksel hizasızlık gibi faktörlerin süreçleri yavaşlattığı ve hata riskini artırdığı belirlenmiştir. Saha verileri kullanılarak geliştirilen simülasyon modeli ile sistem analiz edilmiş ve iyileştirme alanları tespit edilmiştir.

Bu doğrultuda, kontrol noktasının hizalanması, QR tabanlı veri aktarımı ve otomatik belge çıktısı gibi çözümler önerilmiştir. Önerilen sistemin işlem sürelerini azaltarak operasyonel verimliliği artırması ve insan hatasını minimize etmesi beklenmektedir.

Anahtar Sözcükler: Kantar eniyilemesi, simülasyon, süreç iyileştirme, lojistik.

Optimization of Weighbridge Processes

Abstract

This study focuses on inefficiencies in weigh bridge operations at Limak Cement facilities caused by high truck traffic. The current system relies on manual data entry, verbal communication, and a physically misaligned control booth, leading to increased processing times and higher risk of human error. A simulation model was developed based on on-site data to analyze system performance and identify bottlenecks.

Based on the findings, improvements such as control booth alignment, QR-based data transfer, and automated document printing were proposed. The proposed system is expected to reduce processing times, improve operational efficiency, and minimize human errors.

Keywords: Weigh bridge optimization, simulation, process improvement, logistics.

9.1 Company Description

Limak Holding was founded in 1976 in Ankara as a small construction company and has since grown into one of Turkey's leading conglomerates (Limak Group of Companies, 2025). Limak Cement is one of the leading cement producers in Türkiye, operating large-scale facilities with high-volume logistics operations. Due to the continuous flow of inbound raw materials and outbound finished products, a significant number of trucks enter and exit the plant daily, making weigh bridge operations a critical component for maintaining an efficient and uninterrupted logistics flow.

The system consists of two main operational flows: the sales process and the raw material procurement process. In the sales process, trucks arrive empty, are weighed at the entrance, proceed to loading, and then pass through the exit weigh bridge before leaving. In the procurement process, trucks arrive loaded, are weighed upon entry, directed to unloading areas, and exit the system after final weighing.

In both flows, trucks interact with entrance and exit weigh bridges connected to a shared control cabin operated by a single operator. The operator handles both weighing and data entry tasks, creating a tightly coupled system where physical operations and information flow are highly interdependent.

The system also relies heavily on manual and verbal communication. Drivers must leave their vehicles to provide information, which is manually recorded by the operator. This structure increases dependency on human interaction, leading to operational delays and making the system more vulnerable to disruptions.

9.2 Definition of the Engineering Problem

The weigh bridge operations at Limak Anka Cement play a critical role in managing the high volume of inbound and outbound truck traffic. The system consists of two main operational flows: the sales process and the raw material procurement process. In both flows, trucks are required to pass through entrance and exit weigh bridges, while interacting with a shared control cabin operated by a single operator. This structure creates a tightly coupled system where physical operations and information flow are highly interdependent.

Field observations and system analysis revealed that the current system suffers from two major sources of inefficiency: prolonged operation times and human-induced errors. These issues not only affect the throughput of the weigh bridge but also have broader implications on operational reliability and reporting accuracy.

The first problem arises from the operational structure of the weigh bridge process. The interaction between the truck driver and the operator is largely manual and requires the driver to leave the vehicle, communicate necessary information, and return to the truck. These non-value-added activities significantly increase service times at both entrance and exit points. Since both operations are handled by the same operator and infrastructure, even small delays accumulate and result in queue formation, especially during peak hours. As a result, the system experiences unnecessary waiting times and reduced operational efficiency.

The second and more critical issue is human error caused by manual data handling. Information such as company name, order number, material type, and date is verbally communicated by the driver and manually entered into the system by the operator. This process is highly prone to errors including incorrect data entry, missing information, and miscommunication. These errors do not remain local to the weigh bridge operations; instead, they propagate through the system due to SAP integration, directly affecting financial records, reporting accuracy, and operational decision-making.

Moreover, the impact of human error extends beyond the initial mistake. Once an error occurs, it triggers a time-consuming correction process involving multiple stakeholders, manual verification, report revisions, and in some cases communication with external parties. This significantly increases the operational workload and leads to additional hidden costs in terms of time and labor.

In summary, the current weigh bridge system is constrained not by physical capacity limitations, but by process inefficiencies and human-dependent operations. The combination of prolonged service times and error-prone manual data handling creates a system that is both time-inefficient and vulnerable to operational disruptions. Therefore, addressing these two core problems is essential for improving system performance, reliability, and overall efficiency.

9.3 Model Development

To accurately analyze the current operations and evaluate potential improvements, a simulation model was developed based on real system observations and collected data.

For the sales (cement shipment) process, critical time components such as weigh bridge processing times, manual information exchange durations, and loading (filling) times were recorded during on-site visits. These time measurements reflect real operational conditions, including delays caused by drivers exiting their vehicles and interacting with the control booth. Based on these observations, a detailed simulation model Figure 9.1 representing the sales process flow was developed in Arena Simulation Software.

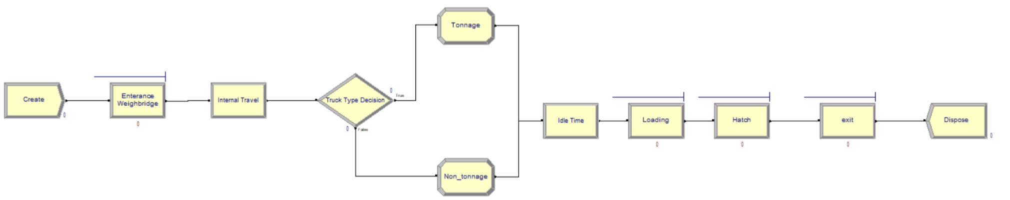

Figure 9.1: Sales Model

For the raw material delivery process, weigh bridge processing times for incoming trucks were recorded. Since there is no operational procedure at the exit weigh bridge for these trucks, no additional processing time was considered for the outbound stage. The process at entry is generally limited to delivery note submission, making it simpler compared to the sales side. However, in some cases, if the truck belongs to Limak itself, delivery note verification may not be required or may be handled in a simpler way. Additionally, factors such as raw material unloading durations and the availability of unloading space were identified and analyzed using historical data. These elements were incorporated into the model to better reflect real system constraints and variability. A separate simulation structure was developed to represent this process.

Using the collected and analyzed data, a discrete-event simulation model Figure 9.2 was developed in Arena Simulation Software. The model represents the full flow of trucks through the system, including arrival, weighing, loading/unloading, and exit processes. Separate logic structures were created to capture the operational differences between raw material and sales processes.

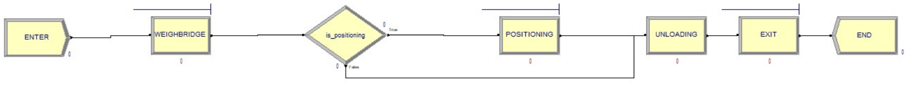

Figure 9.2: Raw Material Model

Although the primary objective of the company was to reduce errors by eliminating manual processes, the Arena simulation also makes it possible to observe the impact of the proposed solutions on system performance, particularly in terms of processing times. The model was developed based on the ANKA factory; however, since similar operational structures exist across different facilities, it can also be used as a prototype for future applications.

Furthermore, the model enables scenario analysis by allowing changes in system parameters. For example, one of the initially considered improvement scenarios was the addition of an extra weigh bridge. Such scenarios can be tested within the simulation environment to evaluate their impact on system throughput and overall performance before real-world implementation.

Overall, the developed model serves as a reliable baseline representation of current operations and provides a solid foundation for testing and evaluating proposed system improvements.

9.4 Validation

To validate the proposed simulation model, real system data provided by the company were compared with the simulation outputs. The validation focused on two key performance measures: the daily number of incoming trucks and the average time spent in the system.

For the sales process, the simulation results closely matched the real system observations. The average number of daily incoming trucks was approximately 104 in the real system and 103 in the simulation. Similarly, the average system time was observed as 38 minutes in both cases, with a 95% confidence interval of $(35.5, 39.5)$. A one-sample t-test was conducted, yielding a test statistic of $t = -0.56$, which is below the critical value $t = 2.262$. Therefore, the null hypothesis could not be rejected, indicating no statistically significant difference between the real system and simulation results.

A similar validation approach was applied to the raw material procurement process. The average system time was approximately 17 minutes in the real system and 17.22 minutes in the simulation, with a 95% confidence interval of $[16.9, 17.4]$. The t-test result $t = 1.414$ was again below the critical value $t = 2.776$, and the p-value 0.23 was greater than 0.05. Hence, the null hypothesis could not be rejected.

Overall, the statistical analysis demonstrates that the simulation outputs are consistent with real system data, confirming that the model provides a reliable representation of the current operations.

9.5 Proposed Solutions

The proposed system is designed to address the identified inefficiencies by reducing manual interaction, eliminating physical misalignments, and improving both operational speed and data accuracy through three integrated improvements targeting operational delays and human errors.

The first improvement focuses on the physical redesign of the control booth to reduce non-value-added operation time. In the current system, the misalignment between booth window and truck driver's window forces drivers to exit their vehicles, creating unnecessary delays. By aligning booth structure with driver's window, communication can be carried out directly without requiring driver to leave the vehicle. This eliminates unnecessary movements, shortens processing time, and improves operational flow.

The second improvement introduces a QR-based digital information transfer system to address human errors caused by manual data entry. Instead of relying on verbal communication, drivers present QR codes containing shipment information, scanned and automatically transferred into the system. This largely eliminates manual data entry, significantly reducing error rates while improving data accuracy, traceability, and process speed.

The third improvement targets delays in the exit process by integrating an industrial printing system for delivery documents. In the current system, delivery notes are manually prepared and handed to drivers, requiring additional interaction and time. With the proposed system, documents generated by the system are automatically printed and positioned in a way that allows drivers to receive them without leaving their vehicles. This reduces service time at the exit stage and minimizes process interruptions.

These solutions are supported by detailed equipment selection, installation planning, and cost analysis to ensure feasibility under real operating conditions. Together, they provide a practical and integrated approach to reducing delays, minimizing human error, and improving overall system performance.

9.6 Benefits to the Company

9.6.1 Human Error Reduction

Based on the Raw Material Unloading Arena simulation results, the primary bottleneck was not identified at the weigh bridge entrance in terms of processing capacity. This indicates that the system is not constrained by physical limitations such as service time or resource availability. Instead, the main source of inefficiency originates from manual operations, particularly human errors occurring during data entry.

These errors, including incorrect date entries, order mismatches, and incorrect monthly record assignments, propagate through the system and directly affect SAP-integrated reports. As a result, incorrect data is reflected in purchasing and operational reports, requiring daily manual verification by employees. Once detected, these errors trigger a multi-step correction process involving incident documentation, system updates, and report revisions. This leads to significant time loss, increased workload, and reduced reliability of operational data.

To address this issue, the proposed integrated automation system minimizes manual data entry and ensures controlled and consistent data processing. By reducing the occurrence of human errors at the source, the system eliminates the need for repetitive correction processes and improves overall operational efficiency. Figure 9.3 presents the elimination of the workload required to solve issues caused by data errors.

Performance Indicator	**Current State**	**Improved State**	**Impact**
Error occurrence	Once–twice per month	Few times per year	Significant reduction
Estimated annual Error Recovery process duration	72 – 120 days	6 – 20 days	**66 – 100 days reduction**
Estimated annual active workload	192 – 288 hours	16 – 48 hours	**144 – 240 hours reduction**
Daily control requirement	Mandatory	Reduced	Lower operational burden
Data reliability	Low	High	Improved decision-making

Figure 9.3: Performance Comparisons

The proposed system reduces the total annual delay caused by human errors by up to 100 days and decreases manual workload by up to 240 hours per year. These improvements significantly enhance data accuracy, reduce dependency on manual intervention, and enable faster and more reliable decision-making processes within the company.

9.6.2 Processing Time Reduction

In addition to eliminating human-induced inefficiencies, the proposed system improves operational performance for the sale operations by reducing processing time at the entrance weigh bridge.

In the current system, the entrance mean is 121.8 seconds with a standard deviation of 17.2, and the exit mean is 91.8 seconds with a standard deviation of 30.2.Therefore, manual data entry, operator–driver interaction, and document handling result in an average processing time of approximately 212 seconds per truck.

With the implementation of the integrated automation system, these manual steps are minimized, and processing becomes faster and more standardized. With the implemented solutions, the entrance and exit processes are significantly improved.

As a result, the total processing time is reduced to approximately 56 seconds per truck.

FOR SALE MODEL;							
Performance Indicator	**Current State**			**Integrated Automaton**			**Impact**
	Minimum	Maksimum	Mean	Minimum	Maksimum	Mean	
Total processing time at the both entrance and exit weighbridges			**~212 seconds**			~56 seconds	~73% reduction
Time spent in the system	~13.8 min	~43.4 min	**~39 min**	**~9.6 min**	**~ 38.9 min**	**~ 35.9 min**	~6–7% improvement
Queue levels	~0 sec	~50 sec	~32 sec	~0 sec	~3 sec	~0.5–1 sec	~96–97% improvement

Figure 9.4: Results of Implementations

Figure 9.4 shows the expected results of the implementation and its benefits. The reduction in processing time leads to faster throughput, shorter waiting times, and improved utilization of weigh bridge resources. This contributes to smoother operational flow within the facility and enhances overall logistics performance.

9.7 Implementation

System components were evaluated in terms of both technical performance and cost, and the detailed cost ranges are presented in Figure 9.5 .

For the QR reading system, industrial devices within the range of 700–2,695 USD were analyzed, and the Datalogic Matrix 220 (Datalogic, 2026) was selected as the most suitable option. The total cost, including installation, was estimated at 1,700 – 2,300 USD per unit.

For the printing system, alternatives between 740–2,850 USD were evaluated, and the Kyocera ECOSYS MA4500ix (Kyocera Document Solutions, 2026) was selected due to its compatibility with the existing A4-based system. The total cost, including installation and protective solutions, ranges between 900 – 1,265 USD.

Protective enclosures were evaluated within the range of 315–1,320 USD for printers and 300–1,500 USD for QR systems. Infrastructure components, including the portable office unit (13,720–15,000 USD) and control cabin improvements (20,580–22,000 USD), were also analyzed.

In addition, SAP integration costs were estimated at 5,720–8,000 USD, while installation and other operational expenses were estimated between 3,430–9,000 USD. When the combined cost of all components were analyzed, the total system installation cost is estimated to be approximately 46,665–60,385 USD. This investment improves system reliability, reduces operational errors, and ensures sustainable long-term performance.

Component	**Quantity**	**Unit Cost (USD)**	**Total Cost (USD)**
Industrial QR Scanner	**2**	**700 – 2,695**	**1,400 – 5,390**
Industrial Thermal Printer	**2**	**740–2,850**	**1480 – 5,700**
Printer Protection Enclosure	**2**	**315 – 1,320**	**630 – 2,640**
QR Protection Enclosure	**2**	**300 – 1,500**	**600 – 3,000**
Portable Office Unit	**1**	**13,720 – 15,000**	**13,720 – 15,000**
Control Cabin Construction	**1**	**20,580 – 22,000**	**20,580 – 22,000**
SAP Integration & IT Setup	–	–	**5,720 – 8,000**
Installation & Miscellaneous Costs	–	–	**820 – 3435**

Figure 9.5: Components and Costs

9.8 Conclusion

This study addressed inefficiencies in weigh bridge operations at Limak Cement by analyzing the current system and identifying key bottlenecks caused by manual processes and physical limitations. A simulation model based on real operational data was developed to represent the system and evaluate its performance.

Based on the analysis, practical and feasible improvements were proposed, including control booth alignment, QR-based data transfer, and automated document printing. These solutions aim to reduce processing times, minimize human error, and improve overall system efficiency.

The results indicate that the proposed system can significantly enhance operational performance while maintaining feasibility for real-world implementation. Overall, this study demonstrates how data-driven analysis and simple technological integrations can lead to substantial improvements in industrial operations.

Bibliography

Datalogic (2026). Matrix 220 industrial scanner. Accessed: April 10, 2026.

Kyocera Document Solutions (2026). Ecosys ma4500ix. Accessed: April 10, 2026.

Limak Group of Companies (2025). Our founders. Accessed: October 21, 2025.

10 Kaynak Atölyesi İş Çizelgelemesinin Eniyilenmesi

Teknopar Endüstriyel Otomasyon

Proje Ekibi

Mina Aktaş, Mehmet Berkay Demirtaş, Ezgi Eker,
Ege Ekmekçi, Arın Küçükgüçlü, Ece Saral, Zeynep Türkoğlu

Şirket Danışmanı
Ertuğrul Bilgiç
Yazılım Mühendisi

Akademik Danışman
Dr. Emre Uzun
Endüstri Mühendisliği Bölümü

Özet

Bu proje, iş atamalarının ve üretim planlamasının elle yapıldığı; bu durumun dengesiz iş yükleri, kaçırılan teslim tarihleri ve atıl süreler gibi verimsizliklere yol açtığı Teknopar Endüstriyel Otomasyon'un kaynak atölyesi için bir Karar Destek Sistemi (KDS) geliştirerek verimliliği artırmayı amaçlamaktadır. Sorun, Karma Tam Sayılı Programlama tabanlı bir matematiksel yapı ile Giffler-Thompson algoritmasını Belirgin Gecikme Maliyeti kuralıyla birleştiren melez bir sezgisel yöntem kullanan Dinamik Esnek Atölye Tipi Çizelgeleme Problemi olarak modellenmiştir. Geliştirilen Python tabanlı KDS, Karma Tam Sayılı Programlama ve sezgisel yaklaşımları birleştiren bir eniyileme motoru ile gerçek zamanlı çizelgeleme ve Gantt çizeneği aracılığıyla görselleştirme sağlayan bir çıktı bölümünü entegre etmektedir. Önerilen çözüm, enyüksek gecikmeyi %3,06 oranında azaltırken istasyon ve makine kullanım oranlarını %11,05 oranında artırmıştır.

Anahtar Sözcükler: Karar Destek Sistemi, Çizelgeleme, Endüstri 4.0, Karma Tam Sayılı Programlama, Esnek Atölye Tipi Çizelgeleme

Welding Process Scheduling Optimization

Abstract

This project aims to improve efficiency by developing a Decision Support System (DSS) for Teknopar Industrial Automations' welding facility, where job assignments and production planning are currently done manually, resulting in inefficiencies such as unbalanced workloads, missed due dates, and idle time. The issue is modeled as a Dynamic Flexible Job Shop Scheduling Problem (DJSSP), which utilizes a mathematical structure, Mixed Integer Programming (MIP), and a hybrid heuristic that combines the Giffler Thompson algorithm with the Apparent Tardiness Cost (ATC) rule. The developed Python-based DSS integrates an optimization engine that combines MIP and heuristics, along with an output section for real-time scheduling and visualization via Gantt charts. The proposed solution has increased station and machine utilization rates by 11.05%, while reducing maximum tardiness by 3.06%.

Keywords: Decision Support System, Scheduling, Industry 4.0, Mixed-Integer Programming, Flexible Job Shop Scheduling.

10.1 Company Description

Teknopar Endüstriyel Otomasyon is a research and development based technology company founded in Ankara, Türkiye, in 1996. Teknopar delivers hardware and software solutions across various sectors such as energy, defense, and mobility. The company operates in four locations and employs over 100 engineers focused on industrial automation and digital transformation. To advance its capabilities in machining and production, Teknopar established a specialized facility in Ankara Aerospace Specialized Organized Industrial Zone (HAB), which serves as a welding and production facility. With its A-Level competence certification from the Ministry of National Defense, the company plays a crucial role in the transition to Industry 4.0 by leveraging data, IoT, and AI to optimize resource utilization and production efficiency.

10.2 System Analysis and Problem Definition

This section evaluates the overall flow of Teknopar's welding facility and identifies core scheduling opportunities that can benefit from a transition to an automated DSS.

10.2.1 System Analysis

The production system at Teknopar's HAB facility follows a multi-station job shop structure. The first stage involves the arrival of pre-cut or bent metal components, which are then transferred to the nine specialized workstations. Six of the stations are designated for small parts while three of them are designed for larger parts. Welding operations are done using twelve mobile machines, including nine gas welders and three specialized argon units (air cooled, water cooled, and AC-DC).

While Teknopar has an integrated Advanced Planning and Scheduling (APS) system for CNC Machines, the real-time data integration does not extend to the welding workshop in HAB. Because of this, the assignments are made manually, relying on the department's experience and assessment of operator availability rather than on data and algorithms.

10.2.2 Problem Definition

The current system faces operational challenges due to its manual scheduling process and frequent disruptions. Relying on historical averages and operator experience often results in extended durations, leading to significant tardiness of the production process. The complexity of precedence relations of operations adds another difficulty to scheduling, and the arrival of high-priority jobs increases scheduling complexity. By transitioning from a manual approach to a DSS, the project aims to stabilize the scheduling process and reduce maximum tardiness, ensuring the facility meets its strict delivery standards required for the defense sector.

10.3 Model Development

10.3.1 Assumptions

To ensure the model functions and is computable, some assumptions were made. The assumptions are as follows:

- Each operation must be completed before its successor can begin.
- Differences in the performance of operators are not modeled explicitly, with an average rate being used for each machine and part combination.
- Machine repair and maintenance times are only taken into account if they are to be scheduled in advance. In the case of unexpected breakdowns, they are handled through rescheduling.
- Transport times between workstations are negligible compared to welding and setup duration.

- Each machine and each station can process at most one operation at a time.
- Machine–station compatibility is predefined, and each operation can only be assigned to feasible pairs.
- The system operates under deterministic processing times and known due dates.

10.3.2 Conceptual Model

As shown in Fig. 10.1, the conceptual model aims to transform operational production data into a feasible and efficient scheduling plan for the welding facility. The main objective of the model is to minimize maximum tardiness while ensuring that all operations are completed within their respective deadlines. To achieve this, the model considers job due dates, the sequence dependencies between operations, and the compatibility of machines and stations for each task. The DSS processes these inputs and constructs a schedule by respecting precedence relations, resource availability, and processing constraints. By using these elements, the model generates a production plan that reflects real operational limitations while improving overall efficiency and delivery performance. The model gives an output consisting of optimized machine and station assignments along with a time-based schedule that minimizes delays and ensures minimum maximum tardiness.

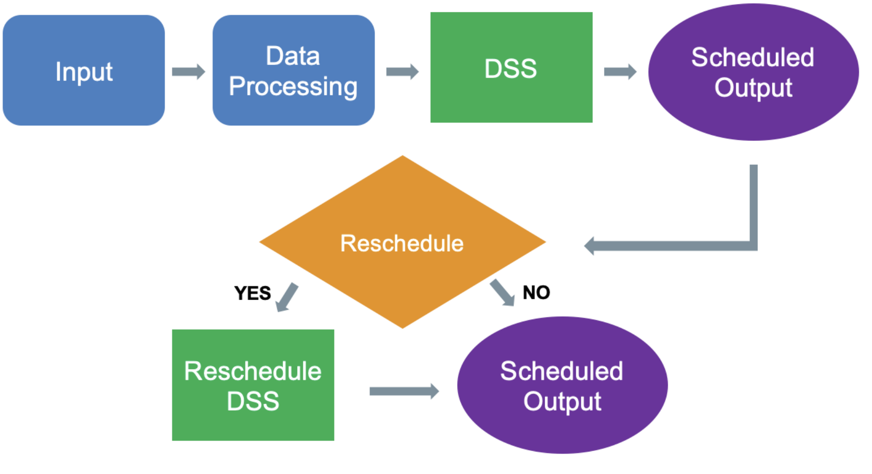

Figure 10.1: Conceptual Model Flow Chart

10.3.3 Mathematical Model

The model is shown in the appendix. The objective function (1) minimizes the maximum tardiness. Constraint (2) ensures that every operation i is uniquely assigned to a machine m. Constraint (3) guarantees that every operation i is assigned to exactly one station l. Constraints (4) and (5) enforce that if operation i is assigned to machine m, then its processing duration $(C_i - S_i)$ equals the processing time P_{im}. Otherwise, the large constant M relaxes the constraint, making it inactive. Constraint (6) prevents large welding operations from being assigned to small stations. Constraint (7) enforces precedence relations : if operation i' precedes operation i, then $S_i \geq C_{i'}$. Constraint (8) ensures that the first operation of job j does not start earlier than its earliest possible start time r_j. Constraints (9) and (10) apply a Big-M formulation to prevent overlapping on shared machines, while Constraint (11) enforces sequencing when operations i and i' are assigned to the same machine. Similarly, Constraints (12) and (13) prevent simultaneous processing on the same station, and Constraint (14) applies sequencing logic only when operations share a station. Constraint (15) sets the welding completion time of job j equal to the latest operation completion time. Constraint (16) computes the completion time C_j^{final} by adding applicable grinding and painting times. Constraint (17) determines tardiness by subtracting the due date from the completion time. Constraints (18) and (19) ensure that the makespan $C_{\max}$ and maximum tardiness exceed all individual completion and tardiness values. Finally, Constraints (20), (21), and (22) impose domain requirements, ensuring binary, non-negative, or integer variable definitions.

10.3.4 Heuristic Algorithm

A heuristic method was also developed to quickly produce viable schedules, as solving the MIP model in large-scale scenarios was computationally inefficient. This approach uses a constructive scheduling process while maintaining the model's essential constraints, such as capacity limits, machine–station compatibility, and precedence relations.

This time-based algorithm identifies tasks whose predecessors are finished at each stage, selecting only those that can proceed given the available machines and stations. When resource conflicts arise between competing operations, the Giffler–Thompson method is used to determine the feasible set of tasks to include in the schedule. To prioritize these operations, the Apparent Tardiness Cost (ATC) rule is used, accounting for deadlines, processing times, and slack. This integration ensures the heuristic can successfully mediate between time constraints and resource optimization. By assigning operations individually and updating resource capacity and timing as needed, the schedule is developed iteratively. This continues until every task is placed, resulting in a valid production plan that defines start and finish times and equipment assignments.

10.4 Validation

Validation was conducted using the projection data from February 2026 to September 2026. Input was taken in JSON format and included information such as precedence relations, duration times, station, and machine type information. To perform validation, certain assumptions were made. It was assumed that all operations were completed within a 09:00–17:00 shift, and operations sharing the same station or machine were distributed evenly across the 8-hour daily production window.

To validate and evaluate our approach, we used several performance indicators. First, machine and station utilization rates were computed from active working times, allowing us to assess how efficiently available capacity was used. Second, total completion time was measured as the maximum finishing time of all operations, providing an indicator of overall production efficiency and time performance. In addition, the workload distribution across machines and stations was analyzed to determine whether the model reduced bottlenecks and ensured balanced task allocation. Idle times were also evaluated to measure unused capacity. To ensure feasibility, all generated schedules were checked against operational constraints. Precedence constraints guaranteed the correct sequence of operations, while resource constraints prevented overlapping assignments on the same machine or station. Finally, the results were compared both numerically, using the defined performance metrics, and visually through Gantt chart representations (Bitran et al., 1982).

The validation results showed a more efficient and reliable schedule. The number of on-time jobs increased from 20 to 27, while tardy jobs decreased from 13 to 6. The makespan decreased from 172.6 to 169.6 hours, and the maximum tardiness decreased from 19.6 to 19 hours. Overall, this combined evaluation framework enabled a systematic comparison of the proposed approach with current operations in terms of efficiency, feasibility, and resource utilization.

10.5 Integration and Implementation

Following validation, the system was integrated into Teknopar's existing production planning environment. We aimed to ensure that the developed solution is not only theoretically sound but also usable within the company's operational structure.

10.5.1 System Integration

The DSS was developed to work alongside Teknopar's current planning processes. It converts the existing data into optimized schedules according to the company's constraints and infrastructure. The system integrates into the existing production planning process by allowing production planners to provide input data such as job details, processing times, and machine availability. This data is then processed by the DSS through the optimization model to generate an efficient production schedule. The resulting schedule is subsequently used to define the final production plan, working in coordination with the company's current system This approach ensures a smooth transition from manual planning to a semi-automated, data-driven system.

10.5.2 Implementation Structure

The implemented system has several features. It includes a user interface that allows production planners to upload data and interact with the system, an optimization engine that runs the mathematical model and heuristic methods, and an output module that generates and exports the resulting production schedules. As can be seen from Fig 10.2, the interface allows users to access the schedule and Gantt charts. Also, users can reschedule for the unavailable machines or stations and reassign the operations manually.

10.5.3 Data Flow and Usage

The DSS operates using structured input data via JSON format, which includes job and operation details, processing and setup times, machine and station information, due dates, and priorities. Once the data is uploaded, the system generates an optimized schedule that can be downloaded and applied directly in the production process.

10.5.4 Rescheduling Integration

One of the most important features of the system is the rescheduling capability. When disruptions occur, users can select the type of disruption (urgent job, machine failure, etc.), input updated information, and generate a revised schedule instantly. This allows the company to adapt quickly to unexpected changes without disrupting the entire production plan. The rescheduling extension in the user interface can be seen in Fig 10.3.

10.5.5 Implementation Plan

The implementation process was carried out in stages: development of the DSS and user interface, testing with company data, pilot study in the production environment, and feedback collection and system improvements.

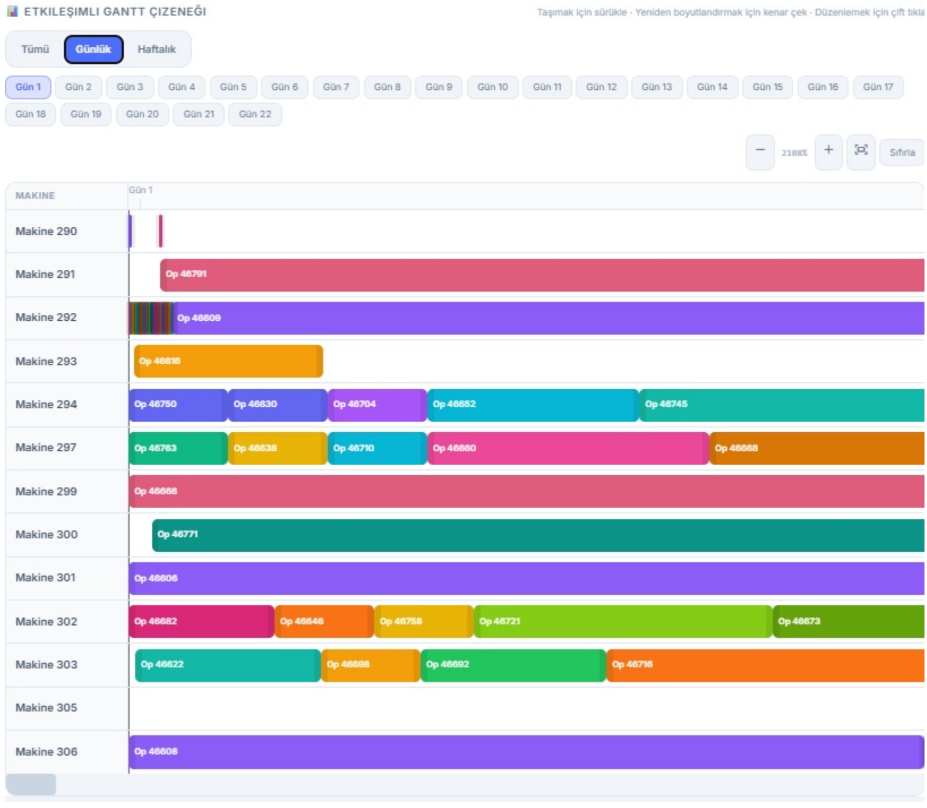

Figure 10.2: User Interface Gantt

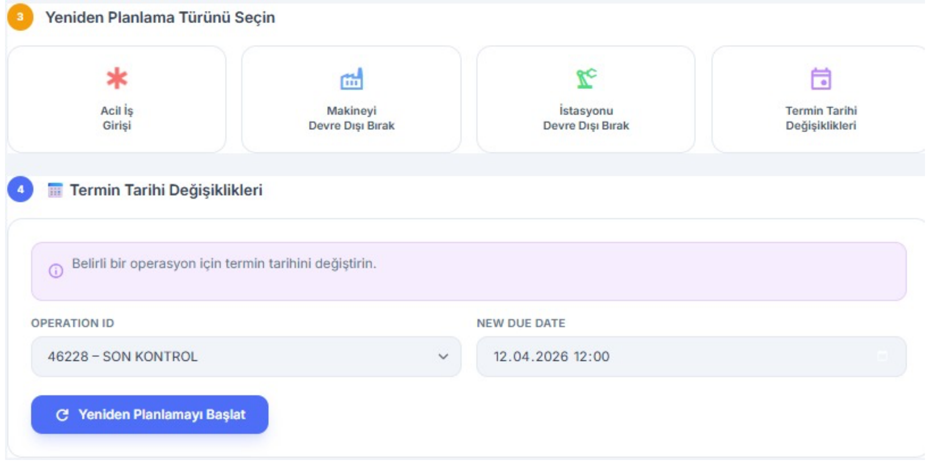

Figure 10.3: Rescheduling User Interface

During the pilot phase, the system was tested under real operating conditions, and its outputs were compared with existing planning methods.

10.5.6 Practical Considerations

The system is designed to be scalable, which is capable of handling large production datasets, flexible, which can adapt to changes in production conditions, and accessible and usable without coding or additional knowledge. To ensure compatibility with Teknopar's operations and usability for planners, the system can facilitate a smooth, efficient implementation.

10.6 Benefits to the Company

The DSS provides Teknopar an operational advantage by optimizing production flow and improving capacity utilization across its welding stations and machines. The DSS outputs daily production Gantt Charts and JSON-formatted production codes to be compatible with the company's current infrastructure. Some of the advantages provided include reducing maximum tardiness and improving machine and station utilization. For benchmarking, 8 months of planning data were used. Improvement rates can be seen in Table 10.1. The system is also adaptable to urgent situations, such as machine or station malfunctions and changes to due dates.

Table 10.1: Improvement Rates

Metric	**Value**
Decrease in Maximum Tardiness	3.06%
Increase in Machine Utilization	11.05%
Increase in Station Utilization	11.05%

10.7 Conclusions

In conclusion, this project develops a data-driven decision-support framework to improve production planning and scheduling at Teknopar's HAB welding facility by modeling the system as a Dynamic Flexible Job Shop Scheduling Problem and implementing a DSS to address inefficiencies arising from manual planning. By integrating mathematical optimization and heuristic methods, the system generates feasible production schedules that improve machine utilization and reduce tardiness while satisfying real operational constraints. In addition, its rescheduling capability enables adaptation to dynamic conditions, such as urgent job arrivals and machine disruptions, supporting faster, more informed decision-making. Future work may include integrating with real-time data sources, enhancing heuristic algorithms, incorporating additional constraints, such as workforce scheduling, and developing improved visualization tools. Overall, the developed system provides a structured solution to support Teknopar's transition toward data-driven production planning and future intelligent manufacturing applications.

Bibliography

Bitran, G. R., E. A. Haas, and A. C. Hax (1982). Hierarchical production planning: A two-stage system. *Operations Research 30*(2), 232–251.

Appendix: Mathemetical Model

1. Sets and Indices

$\mathcal{J} = \{1, \ldots, n\}$: Set of jobs (final products), index $j \in \mathcal{J}$.
$O_j = \{1, \ldots, o_j\}$: Set of bunched welding operations of job j, index $o \in O_j$.
$\mathcal{I} = \{(j, o) : j \in \mathcal{J},\ o \in O_j\}$: Set of all welding operation groups, index $i \equiv (j, o) \in \mathcal{I}$.
$\mathcal{M} = \{1, \ldots, m\}$: Set of machines, index $m \in \mathcal{M}$.
$\mathcal{M}_i \subseteq \mathcal{M}$: Feasible machine set for operation bunch i.
$\mathcal{L} = \{1, \ldots, L\}$: Set of stations, index $\ell \in \mathcal{L}$.
$\mathcal{L}_i \subseteq \mathcal{L}$: Feasible station set for operation bunch i.
$\mathcal{L}_B \subseteq \mathcal{L}$: Set of big stations.
$\mathcal{L}_S \subseteq \mathcal{L}$: Set of small stations.
$\mathrm{Pred}_i \subseteq \mathcal{I}$: Set of immediate predecessors of i.

2. Parameters

$P_{im} \geq 0$: Processing time of operation bunch i on machine m.
$r_j \geq 0$: Earliest possible start time of job j.
$d_j \geq 0$: Due date of job j.
$g_j \in \{0, 1\}$: Grinding requirement of job j.
$p_j \in \{0, 1\}$: Painting requirement of job j.
$t_j^{grind} \geq 0$: Grinding time for job j.
$t_j^{paint} \geq 0$: Painting time for job j.
$\beta_i \in \{0, 1\}$: 1 if operation i must be processed on a big station.
$M \gg 0$: Large constant (machine sequencing big-M).
$M_L \gg 0$: Large constant (station sequencing big-M).

3. Decision Variables

$x_{im} \in \{0, 1\}$: 1 if operation bunch i is processed on machine m.
$y_{i\ell} \in \{0, 1\}$: 1 if operation bunch i is processed at station ℓ.
$z_{ii'm}^{\mathcal{M}} \in \{0, 1\}$: Machine ordering variable.
$z_{ii'\ell}^{L} \in \{0, 1\}$: Station ordering variable.

$S_i \geq 0$: Start time of operation i.
$C_i \geq 0$: Completion time of operation i.
$C_j^{welding} \geq 0$: Welding completion time of job j.
$C_j^{final} \geq 0$: Final completion time for job j.
$T_j \in \mathbb{Q}$: Tardiness of job j.
$C_{max} \geq 0$: Makespan.
$T_{max} \in \mathbb{Q}$: Maximum tardiness.

4. Mathematical Model

$$
\begin{aligned}
&\min T_{max} && && (1)\\
&\text{s.t.} \sum_{m\in\mathcal{M}_i} x_{im} = 1 && i \in \mathcal{I} && (2)\\
&\sum_{\ell\in\mathcal{L}_i} y_{i\ell} = 1 && i \in \mathcal{I} && (3)\\
&C_i - S_i \geq P_{im} - M(1 - x_{im}) && i \in \mathcal{I},\ m \in \mathcal{M}_i && (4)\\
&C_i - S_i \leq P_{im} + M(1 - x_{im}) && i \in \mathcal{I},\ m \in \mathcal{M}_i && (5)\\
&y_{i\ell} \leq 1 - \beta_i && i \in \mathcal{I}, \ell \in \mathcal{L}_S && (6)\\
&S_i \geq C_{i'} && i \in \mathcal{I}, \imath' \in \text{Pred}_i && (7)\\
&S_{(j,1)} \geq r_j && j \in \mathcal{J} && (8)\\
&S_{i'} \geq C_i - M(1 - z^{\mathcal{M}}_{ii'm}) && i, i' \in \mathcal{I} : i < i',\ m \in \mathcal{M} && (9)\\
&S_i \geq C_{i'} - M z^{\mathcal{M}}_{ii'm} && i, i' \in \mathcal{I} : i < i',\ m \in \mathcal{M} && (10)\\
&z^{\mathcal{M}}_{ii'm} \leq x_{im}, \quad z^{\mathcal{M}}_{ii'm} \leq x_{i'm} && i, i' \in \mathcal{I} : i < i',\ m \in \mathcal{M} && (11)\\
&S_{i'} \geq C_i - M_L(1 - z^{L}_{ii'\ell}) && i, i' \in \mathcal{I} : i < i', \ell \in \mathcal{L} && (12)\\
&S_i \geq C_{i'} - M_L z^{L}_{ii'\ell} && i, i' \in \mathcal{I} : i < i', \ell \in \mathcal{L} && (13)\\
&z^{L}_{ii'\ell} \leq y_{i\ell}, \quad z^{L}_{ii'\ell} \leq y_{i'\ell} && i, i' \in \mathcal{I} : i < i', \ell \in \mathcal{L} && (14)\\
&C_j^{welding} = C_{(j,o_j)} && j \in \mathcal{J} && (15)\\
&C_j^{final} = C_j^{welding} + g_j t_j^{grind} + p_j t_j^{paint} && j \in \mathcal{J} && (16)\\
&T_j \geq C_j^{final} - d_j && j \in \mathcal{J} && (17)\\
&C_{max} \geq C_j^{final} && j \in \mathcal{J} && (18)\\
&T_{max} \geq T_j && j \in \mathcal{J} && (19)\\
&x_{im},\ y_{i\ell},\ z^{\mathcal{M}}_{ii'm},\ z^{L}_{ii'\ell} \in \{0, 1\} && i, i' \in \mathcal{I},\ m \in \mathcal{M}, \ell \in \mathcal{L} && (20)\\
&S_i,\ C_i,\ C_j^{welding},\ C_j^{final},\ C_{max} \geq 0 && i \in \mathcal{I}, \jmath \in \mathcal{J} && (21)\\
&T_j,\ T_{max} \in \mathbb{R} && j \in \mathcal{J} && (22)
\end{aligned}
$$

11 | Lokasyon ve Ürün Bazında Kısa Dönem Talep Tahmini

Mavi

Proje Ekibi
Hilal Aydın, Zeynep Ayhan, Elif Balcı, Yiğit Kahveci
Selen Kılıçaslan, Erol Atahan Kutlu, Abide Sıla Yapıcıoğlu

Şirket Danışmanı
Ensar Emirali
Veri Analitiği Yöneticisi

Akademik Danışman
Prof. Dr. Ülkü Gürler
Endüstri Mühendisliği Bölümü

Özet

Bu projenin amacı Mavi için mümkün olan en düşük seviyede, 6-8 haftalık dönem için yüksek doğrulukta sonuçlar üretebilen şirket içi bir talep tahmin modeli geliştirmektir. Mevcut yapıda dış kaynak üzerinden yürütülen tahminleme süreci, talep dinamiklerinin operasyonel kararlara daha esnek ve doğrudan entegrasyonu noktasında bir gelişim alanı sunmaktadır. Bu doğrultuda, satış verileri analiz edilmiş, ürün-mağaza çiftleri kural tabanlı filtreleme ile seçilmiş ve sonrasında kümelendirilmiştir. Bu kümeler talep yapısına göre uygun modeller, Prophet veya XGBoost, ile eşleştirilmiş ve talep tahminleri yapılmıştır. Geliştirilen talep tahminleme sistemi hem tahmin doğruluğunu artırmış hem de Mavi'nin talep planlama ve tedarik zinciri kararlarını şirket içinde daha etkin biçimde yönetmesine katkı sağlamıştır.

Anahtar Sözcükler: Talep Tahmini, Zaman Serisi Analizi, Kümeleme, Veri Analizi, Perakende, Sezonsallık

Short-term Demand Forecasting Based on Location and Product

Abstract

This project aims to develop an in-house demand forecasting model for Mavi that can generate highly accurate forecasts for a 6–8 week horizon at the lowest possible level. In the current setup, the outsourced forecasting process presents an opportunity for improvement in enabling a more flexible and direct integration of demand dynamics into operational decision-making. In response, sales data were analyzed, product-store pairs were selected through a rule-based filtering and the remaining pairs were clustered according to their demand characteristics. These clusters were then matched with appropriate forecasting models, Prophet or XGBoost, and forecasts were generated accordingly. The resulting demand forecasting system improved forecasting accuracy and strengthened Mavi's ability to manage demand planning and supply chain decisions more effectively in-house.

Keywords: Demand Forecasting, Time-Series Analysis, Clustering, Data Analysis, Retail, Seasonality

11.1 Company Description

Mavi is a global lifestyle and apparel company founded in 1991 in Istanbul, Türkiye, with a strong focus on denim products. Over time, the brand has evolved into a full lifestyle company while maintaining denim as its core offering. Mavi operates in 34 countries through approximately 4000 sales points and supports its operations with a strong omnichannel strategy, including both physical stores and digital platforms. The company offers a wide range of products across different categories and targets mainly young, fashion-conscious customers (Mavi, 2025a).

11.2 System Analysis and Problem

In Türkiye, Mavi operates through a broad retail network of 352 retail stores along with 70 dealer stores, over 570 other selling points and through online channels and offers more than 5,000 products (Mavi, 2025b). This high level of product variety is further increased by differences in fabric type, fit, color and design details, together with rapidly changing fashion trends and customer preferences, making demand forecasting particularly challenging. As a result, Mavi has experienced difficulties in generating accurate forecasts, as well as issues related to product availability and fluctuations in purchasing power.

Currently, Mavi relies on a third-party "black box" forecasting system that uses sales, product, stock, and customer data to generate 8-week forecasts for product distribution. Although the system provides relatively accurate forecasts, it lacks transparency and flexibility, as Mavi cannot access or modify the underlying model. This limits the company's ability to improve the model, adapt it to changing conditions, or generate forecasts at different product or category levels. Therefore, Mavi aims to develop an in-house forecasting model that provides full control, transparency, and adaptability.

The new system is expected to generate accurate weekly forecasts for a 6–8 week horizon at the product and store level, while being efficient enough to run weekly and deliver results within 24 hours. The model is built using historical sales and discount data and is designed to capture key demand patterns such as seasonality, trends, and irregular fluctuations through data-driven analysis.

11.3 Proposed Solution Strategy

11.3.1 Assumptions

During the modelling process, several assumptions were made to handle data limitations. Since the nature of discounts is unknown, any recorded discount is assumed to apply to all items in a purchase and all available items are considered discounted during promotional seasons. As stock information was not disclosed to us, it is assumed that observed sales quantities reflect total realized customer demand and no lost sales occurred due to stockouts. Additionally, customer information was also not shared with us due to confidentiality. Product attributes, such as category and subcategory types, are treated as fixed variables.

11.3.2 Constraints

The geographical scope of the project is restricted to two regions and the scope of the analysis is restricted to three product categories. The available historical data cover the years 2023-2025. The forecasting model is required to generate results within a 24-hour time frame, as the company wants to calculate the demand forecast within a day. For confidentiality reasons, the dataset shared with us was anonymized and some numerical values were altered before analysis.

11.3.3 Objective

The objective of this project is to forecast short-term (6–8 weeks) weekly demand at the product–store pair level, representing the lowest level of aggregation possible, across three product categories and two regions within

24 hours by analyzing and identifying the dynamics that affect demand.

11.3.4 Solution Approach

Our proposed solution is a pipeline that first filters and clusters product-store pairs, and then generates forecasts for each pair using the model assigned to its cluster. The Figure 11.1 shows the diagram describing the proposed pipeline as follows:

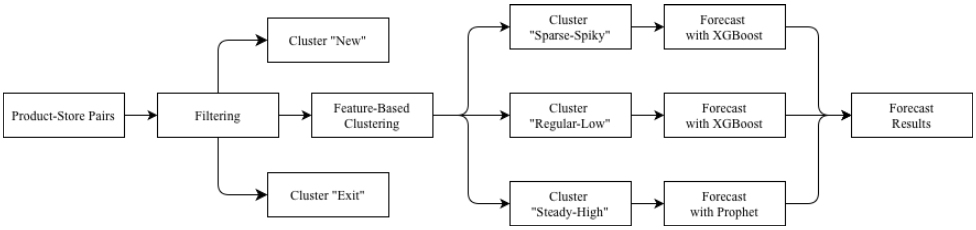

Figure 11.1: Diagram of the Proposed Solution Pipeline

Clustering was introduced because the analysis showed that products with different demand characteristics require different model configurations and, in some cases, entirely different forecasting models. Therefore, to enable more appropriate model selection and high accuracy, product-store pairs with similar demand behaviour are grouped before model training.

Before clustering, history and activity filters were imposed to ensure that the retained product-account pairs were suitable for modelling and forecasting. The pairs excluded through the preprocessing were grouped into two rule-based clusters, New and Exit, and will not be forecasted.

For the remaining pairs, clustering was performed using a hierarchical two-stage K-Means methodology on product-account level sales data. Each series was transformed into a fixed-dimensional feature vector capturing key demand characteristics such as level, variability, intermittency, trend, seasonality, and holiday responsiveness.

In the first stage, top-level K-Means clustering was applied to identify broad structural groups in the data. Cluster validity was evaluated using multiple diagnostics, including internal validation indices, visualisations and cluster feature profiles. In the second stage, the larger top-level cluster was further sub-clustered to test whether additional meaningful structure existed within it. Overall, this resulted in three forecastable clusters, alongside Clusters New and Exit for excluded product-account pairs.

The final clusters are:

- Cluster New: Pairs only sold within the last year, excluded from forecasting

- Cluster Exit: Inactive or discontinued pairs, excluded from forecasting

- Cluster Sparse-Spiky: Pairs with a highly intermittent demand pattern, characterized by many zero-sales periods and occasional sharp spikes

- Cluster Regular-Low: Pairs with a relatively regular demand pattern but consistently low sales volume

- Cluster Steady-High: Pairs with a stable and continuous demand pattern and comparatively high sales volume

The models chosen for the forecasting pipeline are Prophet and XGBoost. Prophet decomposes the time series into trend, seasonality, and holiday effects. It can also handle missing observations and is particularly suitable for series with stable patterns, strong seasonal structure and sufficient historical data. (Phalachandra et al., 2023). Therefore, it is assigned to Cluster Steady-High.

XGBoost utilises ensemble decision trees and captures temporal dynamics through lagged sales features, rolling statistics and calendar-based features. It can learn complex and nonlinear demand patterns directly from data without assuming a fixed trend or seasonal structure and model asymmetric and delayed effects of events (Massaro et al., 2021; Chen and Guestrin, 2016). Since we observed that XGBoost performs better for irregular and less-structured demand behavior, it is assigned to Clusters Sparse-Spiky and Regular-Low. This required the development of two separate XGBoost configurations since the optimal parameters differ according to cluster behavior.

11.4 Validation

To ensure the robustness and reliability of the forecasting models, a comprehensive validation process against real-world data was conducted. Historical sales data provided by the company is the basis for weekly demand forecasts, therefore this step includes comparing the proposed models' forecasts with realized sales using appropriate metrics such as RMSE.

A representative subset of product-store pairs was selected to assess whether the model outputs were consistent in different settings and reflected the outcomes observed in the real system. We checked that the main factors affecting demands, such as seasonality and holiday effects, were captured by the models correctly and their impacts on forecasts were as expected. 2023 and 2024 sales of the pairs were used to train the models and then their forecasts were tested on the 2025 sales.

Each model was updated using hyperparameter tuning through grid search. Comparing the RMSE values of these updated models, we tried to achieve the model that results in the lowest RMSE value for that product-store pair. We selected final models prioritizing RMSE and sMAPE to ensure better overall fit and stability. RMSE is useful for comparing forecast accuracy of different models for a single product-store pair and sMAPE allows for scale-independent evaluation of the models across different pairs. Through hyperparameter tuning, the test sMAPE of a model reduced from 156.30% to 80.67% over a scale of 200% while the baseline seasonal naive model's test sMAPE was found to be 94.33% for a representative pair.

In addition, we did cross-validation to evaluate the models' stability across different time windows and performed residual error analysis to identify potential systematic bias and determine whether the prediction errors are randomly distributed as expected.

The validation results were found to be satisfactory and suggested that the models are reliable and can produce accurate results for different pairs under various conditions.

11.5 Integration and Implementation

At the delivery stage, the proposed system is implemented through a simple and structured pipeline designed to integrate seamlessly into Mavi's existing data and decision-making infrastructure while requiring minimal changes to current operations. The forecasting process follows a structured workflow consisting of data extraction, preprocessing, model execution and forecast generation.

The user provides the required input files, including sales and cluster information, after which the forecasting pipeline is executed automatically. The corresponding forecasting model is then selected and applied according to the cluster assigned to each pair and the resulting weekly demand forecasts are exported in Excel format. The developed interface shown in Figures 11.2 and 11.3 helps standardize the execution procedure, reduce manual intervention and ensure consistency in model usage across different users. The outputs are generated in a format that can be directly integrated into company processes such as inventory planning or store management.

The clustering step of the pipeline is executed only once initially to ensure operational efficiency. When new data becomes available, new product-store pairs are first checked against the New and Exit rules. If neither rule applies, they are assigned to the most appropriate existing cluster based on the distance between their feature values and the cluster centroids, rather than by re-running the clustering algorithm on the entire dataset. It is recommended that cluster definitions and assignments be periodically reviewed and updated to ensure that they remain representative over time. The results showed that these conditions were met and that the system performed successfully in practice.

Finally, the modular design of the pipeline supports scalability across a large number of product-store combinations and allows for straightforward maintenance and future extensions, such as the inclusion of additional regressors or alternative models. Once deployed, model performance can be continuously monitored using established metrics such as RMSE and sMAPE, while a seasonal naive model may be retained as a fallback benchmark to ensure robustness in operational settings.

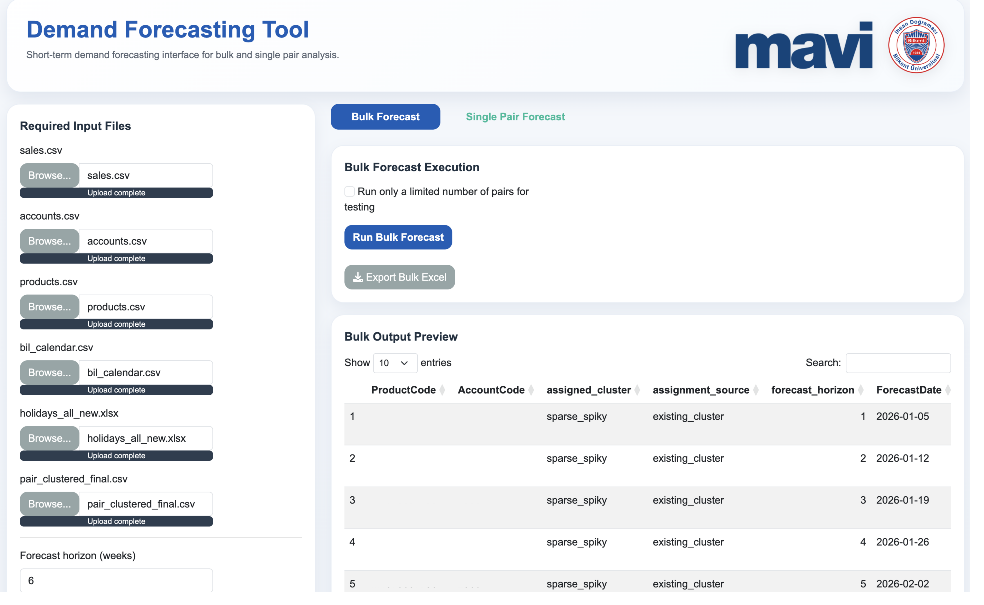

Figure 11.2: User Interface: Bulk Forecasting Page

11.6 Benefits to the Company

To evaluate the effectiveness of the proposed solution approach, our forecasting pipeline was benchmarked against the black-box system currently used by the company. The comparison showed that our approach improved sMAPE by 28.5%, despite generating forecasts at the significantly more granular product-account pair level, whereas the existing system operates only at the product level. Achieving better forecasting performance even at this lower level of aggregation, where demand patterns are inherently more volatile and difficult to predict, is a strong indication of the robustness and practical value of the proposed approach.

In addition to the overall forecasting accuracy, we also evaluated the weekly overestimation and underestimation behavior of the models. Comparing the weekly under and overestimation error levels of our system and the current blackbox, we increased the weekly accuracy by approximately

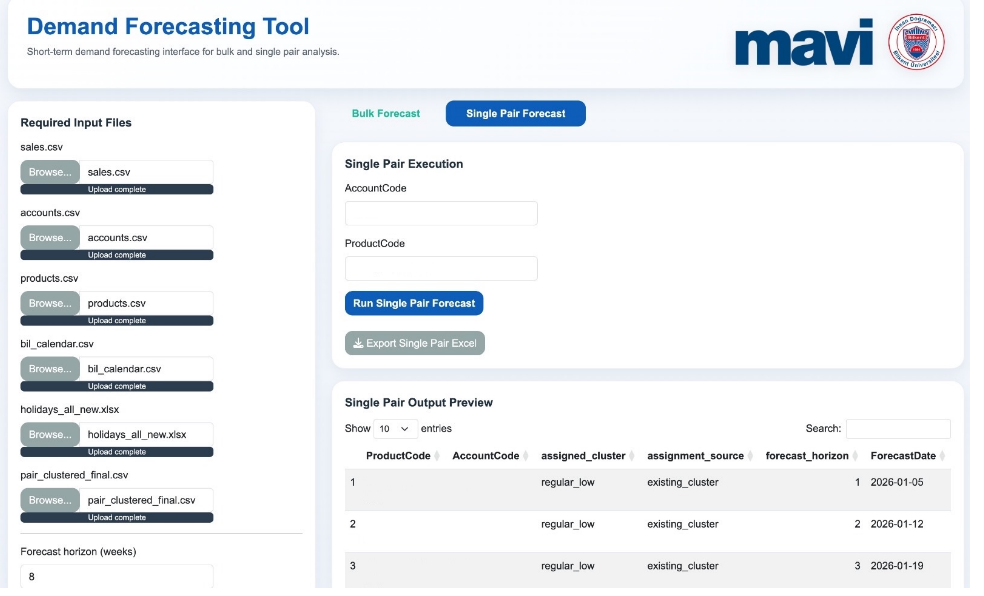

Figure 11.3: User Interface: Single Pair Forecasting Page

40.9%. This gain is not only statistically meaningful but also operationally critical. The improvement at the weekly level indicates that the proposed model captures demand patterns more effectively, supporting more reliable planning decisions. As a result, the proposed solution contributes to more efficient inventory management and improved product availability in stores.

Another important advantage is that the models can be applied flexibly across different product-store segments. The proposed forecasting system can be implemented seamlessly for over 5000 different product types in Mavi's catalogue that are sold at over 4000 selling points in 34 countries. By providing historical sales data of desired product-store pairs and accurate local holiday dates, the system can be adapted to different needs with minimal alterations.

11.7 Conclusion

This project demonstrates that a data-driven forecasting approach can provide strong demand prediction performance in a complex retail environment. By incorporating time series dynamics, seasonality and external factors, the proposed forecasting pipeline provides a scalable and interpretable solution that aligns well with the company's operational requirements. The results indicate that the system meets the organization's expectations in terms of both forecasting accuracy and practical applicability. Beyond forecasting performance, the proposed solution also offers operational and strategic value by reducing dependence on third-party forecasting providers, increasing transparency in the forecasting process and enabling demand planning decisions to be managed more effectively in-house. Improved forecasting performance also helps reduce exposure to stockout risk and excess inventory by supporting more informed replenishment and allocation decisions.

Bibliography

Chen, T. and C. Guestrin (2016). Xgboost: A scalable tree boosting system. *CoRR abs/1603.02754*.

Massaro, A., A. Panarese, D. Giannone, and A. Galiano (2021, 08). Augmented data and xgboost improvement for sales forecasting in the large-scale retail sector. *Applied Sciences 11*, 7793.

Mavi (2025a). Mavi faaliyet raporu 2024. https://www.mavicompany.com/i/assets/documents/2025/MAVIFAALIYETRAPORU2024-2.pdf. Accessed: October 26, 2025.

Mavi (2025b). Mavi global multi channel network. https://www.mavicompany.com/en/mavi-company/global-multi-channel-network. Accessed: October 26, 2025.

Phalachandra, P., P. S. Nukal, R. Ranganath, R. Naik, and A. V. Singh (2023, Aug). Evaluating the predictive power of prophet model and lstm model in retail sales forecasting. In *2023 IEEE 4th Annual Flagship India Council International Subsections Conference (INDISCON)*, pp. 1–7.

Mağaza Yerleşimini Dikkate Alan Depo Ürün Yerleştirme Eniyilemesi için Ardışık Bir Çerçeve

12

A101 Yeni Mağazacılık

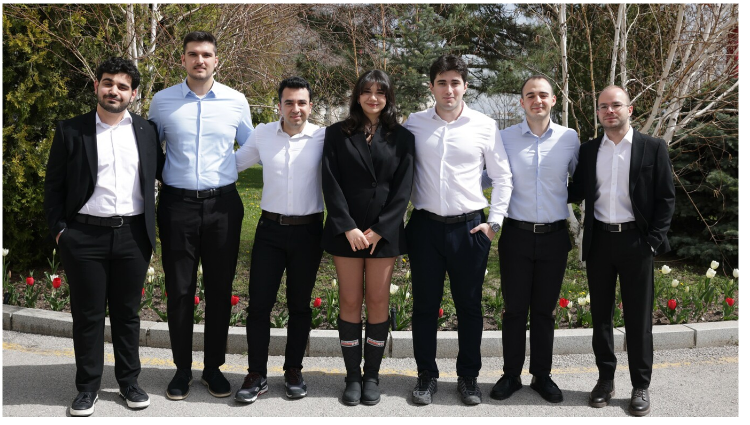

Proje Ekibi

Oğuz Oral, Kaan Öztürk, Erberk Saydan,
Hasan Utku Şenkaya, Kerem Arda Tantekin, Selin Uysal, Oğulcan Yavuz

Şirket Danışmanı
İsmail Gökhan Kaya
Depo Operasyonları ve Projeler Müdürü

Akademik Danışman
Dr. Öğr. Üyesi Gizem Özbaygın
Endüstri Mühendisliği Bölümü

Özet

Bu proje, A101'in depo operasyonlarını veriye dayalı, çok aşamalı bir optimizasyon çerçevesiyle iyileştirmeyi amaçlamaktadır. Mevcut sistemde ürün yerleşimleri deneyime dayalı olarak belirlenmekte, bu da toplayıcıların uzun mesafeler kat etmesine ve verimsiz palet yapılarına yol açmaktadır. Önerilen çerçeve üç aşamadan oluşmaktadır: yüksek etkili ürünleri giriş noktasına yakın konumlandıran Altın Bölge Modeli, tüm ürünleri dayanıklılık sınıflarına göre depo hücrelerine atayan Yerleştirme Modeli ve depo düzenini mağaza raf planıyla uyumlu hale getiren Son İşlem aşaması. Gerçek operasyon verileriyle yapılan performans testleri, palet bazlı toplama mesafesinde ortalama %21,20 iyileşme sağlamıştır. Geliştirilen Karar Destek Sistemi A101'in 58 bölgesel deposuna uyarlanabilir yapıdadır ve şirket yetkililerince gerçekleştirilen fizibilite analizi sonucunda uygulamaya uygun bulunmuştur.

Anahtar Sözcükler: Depo optimizasyonu, ürün yerleştirme, sipariş toplama, paletleme, karar destek sistemi.

A Sequential Framework for Store Layout-Aware Warehouse Slotting Optimization

Abstract

This project optimizes warehouse operations for A101 by replacing experience-based product placement with a data-driven, multi-stage optimization framework. The current system leads to excessive picker travel distances and unstable pallet structures. The proposed framework consists of three phases: a Golden Zone Model that consolidates frequently picked and strongly related products near entry points; a Slotting Model that assigns all products to storage cells based on durability classes and picking frequency; and a stability-constrained post-processing phase that aligns the warehouse layout with downstream store shelf organization through controlled local swaps. Performance testing with real operational data using a pallet-level travel distance simulation showed that the proposed layout achieved an average improvement of 21.20% in picking distance under a controlled single-pallet benchmark agreed with the company as the most operationally representative measure. The developed Decision Support System is scalable across A101's 58 regional warehouses, and a feasibility analysis conducted by the company confirmed the solution's readiness for deployment.

Keywords: Warehouse optimization, order picking, slotting, palletizing, decision support system.

12.1 A101 and Problem Identification

12.1.1 Company Description

A101 Yeni Mağazacılık A.Ş. is one of Turkey's largest discount supermarket chains, operating over 13,000 stores across all 81 provinces with more than 70,000 employees (A101, 2025). The company manages 58 regional warehouses, approximately 7,000 warehouse staff, and a central distribution network dispatching daily shipments to every store.

Each warehouse is organized into five temperature zones: Dry storage, Cold (+4°C), Fruit/Vegetable (+12°C), Frozen (−18°C), and Meat/Poultry (0–2°C). The dry storage area uses lettered rows and numbered columns to systematically locate products (Figure 12.1). Selectors navigate the warehouse using handheld terminals, picking items sequentially; each selector currently picks an average of 14 pallets per shift.

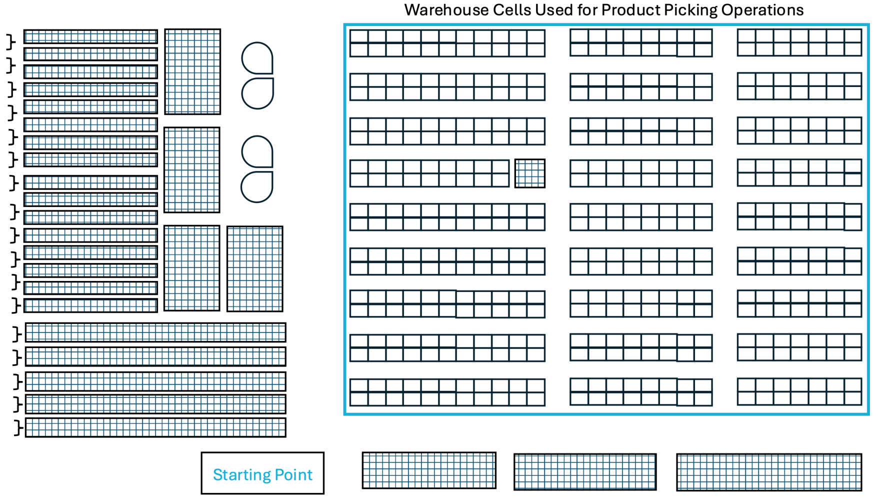

Figure 12.1: Schematic layout of the Esenboğa dry-storage warehouse with lettered rows and numbered columns used for picking operations.

12.1.2 Problem Definition

Analysis of five months of order data from the Esenboğa Warehouse, covering 1,389 products, 17 product groups, 309 goods categories, and 415 orders from 10 stores, revealed significant inefficiencies. A heatmap of storage-location visit frequencies (Figure 12.2) showed that frequently picked products are scattered across distant aisles rather than being spatially consolidated.

Figure 12.2: Heatmap of cell visit frequencies (log scale); darker red cells indicate higher visit counts, green cells are rarely visited.

The root causes are threefold:

- Product placement relies on manager experience rather than data, creating inconsistency across warehouses.
- The palletization order disregards store shelf layouts, increasing shelf-stocking time at the destination store.
- Mixing food and cleaning products on the same pallet violates hygiene regulations and complicates store receiving.

The project's performance measures are total picker walking distance, pallet stability, and alignment with store shelf layouts. The primary deliverable is a Decision Support System (DSS) that generates optimized layouts applicable to all 58 A101 warehouses.

12.2 Proposed Solution Strategy

The proposed framework addresses warehouse slotting through a hierarchical, three-stage optimization pipeline. The warehouse is first partitioned into two operationally independent halves (Food and Cleaning) based on the hard constraint that these categories must not share pallets. Within each half, a three-stage optimization is applied. Recent literature confirms the effectiveness of such coordinated warehouse design approaches (Van Gils et al., 2018; Duque-Jaramillo et al., 2024).

12.2.1 Phase I: Golden Zone Model

The Golden Zone Model selects a focused subset of products for placement in a contiguous, easily accessible region near the warehouse entry point. The model maximizes a weighted combination of individual product frequency and pairwise product affinity, formulated as a Quadratic Knapsack Problem (Gallo et al., 1980). The affinity matrix is constructed from historical joint-purchase patterns (Chan and Chan, 2011), with sparsity filtering to reduce computational complexity. The complete mathematical formulation is given in Appendix 12.A. The capacity parameter is determined heuristically based on the physical layout of corridors connecting fixed start and end points within each warehouse half, resulting in four product clusters: Golden Zone Food, Regular Food, Golden Zone Cleaning, and Regular Cleaning.

12.2.2 Phase II: Slotting Model

The Slotting Model assigns each product to an exact storage location within its designated zone, following a Generalized Assignment Problem architecture (Viveros et al., 2021). The objective minimizes total weighted Manhattan distance between entry/exit points and each product's assigned cell, weighted by picking frequency. A detailed formulation is reported in Appendix 12.B. The Class Zone Matching constraint ensures that heavier products are picked before lighter ones, improving pallet stability during transport. The durability class assignments used by this logic are summarized in Figure 12.3.

1	2	3		4	5
MADEN SUYU SADE & AROMALI	SÜT	ŞEKER	AĞIZ BAKIM	ÇİKOLATALAR	EKMEK
GAZLI İÇECEK KOLA & ZERO&LIGHT	UNLAR	SOĞUK KAHVE	RENKLİ KOZMETİK	GOFRETLER	TUVALET KAĞIDI
ÇAMAŞIR SUYU NORMAL & ULTRA	TURŞULAR	LİMONATA	VÜCUT BAKIM	DEMLİK VE BARDAK POŞET ÇAY	KAĞIT HAVLU
GAZLI İÇECEK MEYVELİ	BİTKİSEL TOHUM YAĞLARI	HELVA	MAYA	BEBEK GIDALARI	YUMURTA
GAZLI İÇECEK GAZOZ	ENERJİ İÇECEĞİ & SPORCU İÇECEĞİ	SOSLAR	KAHVE BEYAZLATICI	ŞEKERLEMELER	CİPSLER
SALÇA	BUZLU ÇAY	HAZIR YEMEK VE MEZELER	BULYON	HİJYENİK PED ULTRA& NORMAL&GÜNLÜK	KAĞIT MENDİL
ZEYTİN	SAÇ BAKIM	KONSERVE	HAYVAN MAMALARI	KRAKERLER	INSTANT KAHVE&3'Ü1&2'Sİ1 ARADA&CAPPUCINO
TUZ	SABUN	KAHVALTILIK SOS-SET	KOKU	BİSKÜVİLER	ÇOCUK BEZİ
YUMUŞATICI NORMAL & KONSANTRE	BAKLİYAT	ÇÖP POŞETİ	KURUYEMİŞ	TEMİZLİK BEZİ&SÜNGERİ&ELDİVENİ&TELİ&FIRÇ	PEÇETE BEYAZ & DESENLİ
SIVI BULAŞIK DETERJANI	SİRKE	EZMELER	MAKARNALAR	TAHIL PATLAĞI	GEVREKLER
SIVI ÇAMAŞIR DETERJANI	REÇEL	TOZ İÇECEKLER	ERİŞTE	TÜRK KAHVESİ	KURU MEYVE
YÜZEY & AHŞAP & CAM & HALI TEMİZLEYİCİ	ŞALGAM SUYU	İRMİK	REÇETESİZ SAĞLIK MAMULLERİ	SAKIZLAR	HAZIR UNLU MAMULLER
MUTFAK&BANYO TEMİZLEYİCİ&LAVABO AÇICI	PİRİNÇ	KEKLER	PASTA MALZEMELERİ	KURABİYE	
LEKE ÇIKARICI TOZ & SIVI	TAHİN-PEKMEZ	TOZ DETERJAN BEYAZ & RENKLİ&ELDE YIKAMA	BİTKİ VE MEYVE ÇAYLARI	İÇ GİYİM	
MANGAL MALZEMELERİ	BULAŞIK MAK. TUZU & PARLATICISI&KOKU GİD	ISLAK HAVLU & ISLAK MENDİL	TRAŞ ÜRÜNLERİ	HAZIR TOZ TATLILAR	
	BULAŞIK MAK.DETERJAN&TABLET & KAPSÜL&JEL	M.NEKTARI-MEYVELİ-ARO. İÇECEK-M.SUYU%100	GRANÜL KAHVE		
	KOLONYALAR	YÜZ BAKIM	ÇORAPLAR		
	BAL	KREMALAR	KULLAN AT ÜRÜNLER		
	TÜL BEYAZLATICI&KİREÇ ÖNLEYİCİ&HALI TEM.	DÖKME ÇAYLAR	BAHARATLAR		
	WC & BANYO KOKU GİDERİCİ	DEO-EDT-ROLL ON/STICK	ÇORBA		
		KREMLER VE LOSYONLAR	PAMUK VE KULAK ÇUBUĞU		

Figure 12.3: Durability class coefficients by product group, ranging from class 1 (heaviest, picked first) to class 5 (lightest, picked last) to ensure pallet stability.

Zone reservations use a dynamic "Flow Ratio" heuristic that assigns each storage slot a normalized flow score based on its relative position between entry and exit points.

12.2.3 Phase III: Store Layout Alignment

After Phase II produces an optimized base layout, a post-processing phase adjusts product positions to improve compatibility with downstream A101 store shelf organization. As illustrated in the representative store layout in Figure 12.4, related product categories occupy contiguous shelf blocks within each A101 store; when products destined for neighboring shelves are picked from distant warehouse locations, store replenishment becomes slower and less efficient. The goal of Phase III is to replicate this spatial coherence within the warehouse itself.

Phase III employs a stability-constrained tie-break swap approach. The Phase II solution is preserved as the base layout; Phase III evaluates pairwise swap candidates and accepts a swap only if four conditions are simultaneously satisfied:

- Same durability class: both products belong to the same class so pallet stability is preserved.
- Close picking frequency: the absolute frequency difference does not exceed a tie-break threshold ($\varepsilon = 0.002$), ensuring swaps occur only among near-identical-priority products.

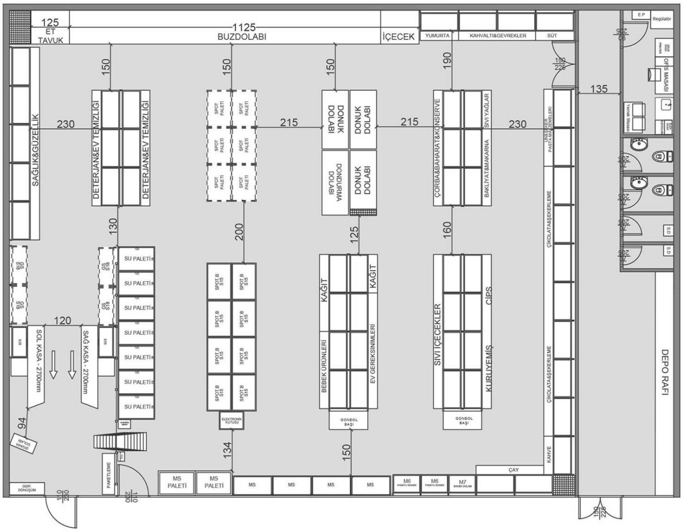

Figure 12.4: Representative A101 store layout with contiguous shelf blocks for related categories. Phase III replicates this spatial coherence in the warehouse.

- Positive family distance gain: the swap reduces total Manhattan distance among products in the same store shelf family ($\Delta > 0$).
- Bounded operational cost: the picking cost increase (frequency $\times$ distance) stays below a tolerance ($\tau = 0.05$).

A hard budget limits affected cells to at most 50. This limit was jointly determined with the industrial advisor specifically to bound the potential negative impact of post-processing on the base optimization; the company required that layout changes remain small enough to be operationally manageable and visually auditable. Each swap affects exactly two cells, yielding at most 25 accepted swaps. At each iteration the algorithm selects the globally best qualifying swap, applies it, and locks the involved products to prevent repeated movement. The process continues until the budget is exhausted or no further qualifying candidates remain, guaranteeing that the selected swaps represent the best possible 50 cell changes under the given constraints.

12.3 Verification, Validation, and Feasibility

Model verification was conducted through controlled tests for both optimization models:

- Slotting Model: Continuity tests confirmed layout stability under uniform frequency scaling, degeneracy tests verified cost sensitivity by penalizing optimal cells, and a mirror-world experiment with inverted distance matrices confirmed correct assignment logic.

- Golden Zone Model: Capacity sensitivity analysis showed monotonic objective improvement with increasing zone size, and dominance experiments confirmed the intended priority structure when frequency or affinity was neutralized.

Validation assessed compatibility with actual A101 operations: model assumptions were checked against physical conditions at the Esenboğa warehouse, and the optimized layout was compared with the current system using historical order sets through a pallet-level travel distance simulation. Additional scenario analyses covered peak demand, low demand, mixed category, fragile-only, and heavy-only orders.

A feasibility analysis was conducted in collaboration with A101's warehouse operations management. The industrial advisor reviewed the DSS outputs, including cell assignments, zone boundaries, and product group mappings, against physical constraints such as aisle widths, shelf capacities, and temperature zone boundaries. The industrial advisor formally approved the solution's applicability, confirming that the proposed layout is operationally deployable under real warehouse conditions.

12.4 Results

Performance was evaluated through a pallet-level picking simulation comparing travel distances under the current and proposed configurations. The DSS replays historical orders and measures total distance per pallet. Durability constraints were respected throughout, and all reported improvement rates represent averages across the full set of test orders. Three comparison settings were used:

- Setting 1, Actual picker order vs. proposed layout with nearest neighbor (NN): This primary comparison measures distance in the current layout using the actual historical selector route against the proposed layout with NN routing. The NN algorithm closely approximates real picking behavior as selectors tend to move toward the next closest feasible location. Average improvement: 61.24%.

- Setting 2, NN on both layouts: Applying the same NN routing to both configurations isolates the pure layout effect by eliminating any influence of suboptimal selector routing in the baseline. Average improvement: 15.06%.

- Setting 3, Single giant pallet: Each store order is treated as one consolidated pallet, removing pallet-allocation confounds and providing a more controlled benchmark. Average improvement: 21.20%.

Following a joint review with the industrial advisor, Setting 3 was jointly identified as the most operationally representative benchmark because it eliminates confounding factors related to pallet allocation and routing variability. The 21.20% improvement was formally agreed upon with the industrial advisor as the reference rate for all subsequent operational benefit projections.

Golden Zone utilization analysis showed that on average 63% of products in an order are picked from the Golden Zone, confirming the substantial contribution of Phase I to the observed distance reduction.

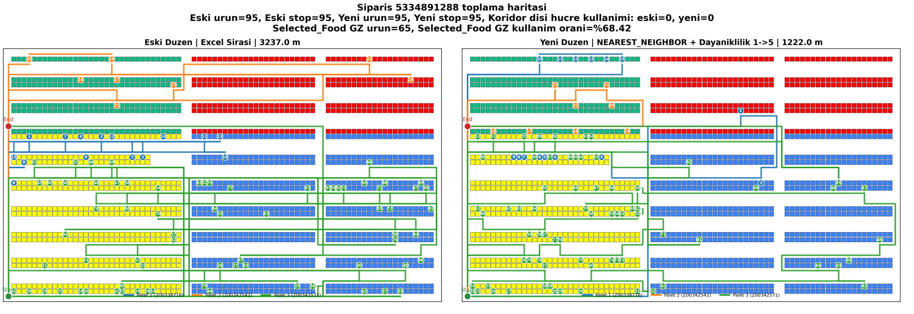

Figure 12.5: DSS route comparison between the current layout (left, 3237.0 m) and the proposed layout (right, 1222.0 m) for the same order. Red and blue lines trace the picker's walking path through warehouse aisles.

Phase III post-processing accepted 25 swaps affecting 50 cells (the full budget) with an overall impact rate of 3.91%. All 14 affected Golden Zone cells fell in the Food section; the Cleaning section remained unchanged. As shown in Figure 12.6, the changes are spatially localized, predominantly short-range movements within the same zone, validating the design intent of controlled, local corrections that improve store-family proximity without disrupting the base solution.

12.5 Integration, Implementation, Benefits

12.5.1 Decision Support System

All optimization outputs are integrated into a scalable DSS whose five-step workflow is shown in Figure 12.7: (1) Data upload via Excel files containing product and warehouse information; (2) Golden Zone analysis execution with configurable parameters; (3) Slotting and cell assignment

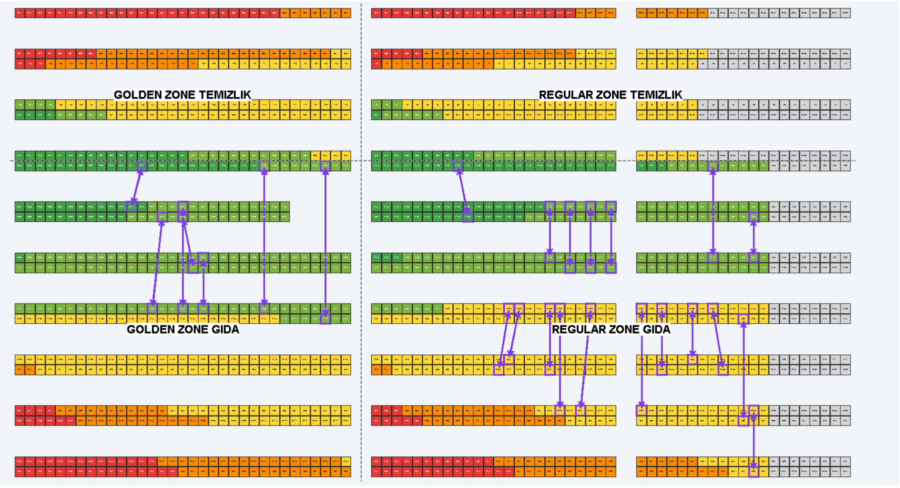

Figure 12.6: Phase III impact map for Esenboğa warehouse. Purple-bordered cells indicate changed product assignments; arrows show movement direction from new to previous cell.

with durability-class enforcement; (4) Warehouse map visualization showing product locations with durability-class color coding; and (5) Performance testing with pallet-level route comparison between current and proposed layouts. The system accepts any warehouse layout as input, making it adaptable to all 58 A101 regional warehouses without requiring model modifications. Warehouse managers can rerun analyses for different warehouse geometries, campaign periods, and product mixes without rebuilding the optimization models, supporting future rollouts and periodic slotting adjustments across the distribution network.

12.5.2 Pilot Study

The Esenboğa warehouse was selected as the pilot site. Following feasibility approval from the industrial advisor, the deployment was scheduled to start in the first week of July 2026, immediately after Eid al-Adha, which marks a lower-demand window suitable for warehouse reorganization. The transition plan requires six Sundays for physical relocation of products, a short adaptation period for selectors to become familiar with the new positions, and a phased feasibility check before broader rollout across other regional warehouses. This confirmed schedule demonstrates that the project has moved beyond the analytical stage into an active implementation pipeline.

12.5.3 Benefits to the Company

The proposed system replaces A101's experience-based placement with a repeatable, data-driven decision process. The durability-based picking logic

Figure 12.7: DSS five-step workflow: data upload (Step 1), Golden Zone analysis (Step 2), slotting (Step 3), map visualization (Step 4), and performance testing (Step 5).

improves pallet stability, strict food–cleaning separation supports hygiene compliance, and Phase III groups products compatibly with store shelf organization, helping store personnel replenish shelves faster.

A key indicator of the framework's effectiveness is the Golden Zone utilization rate: on average, 63% of products in any given order are picked from the Golden Zone, confirming that Phase I successfully consolidates the most frequently needed items into the accessible entry-point region, one of the project's primary design goals.

Using the 21.20% distance improvement as the basis for projections at company scale:

- Increased store coverage capacity: The current workload is completed with only 78.8% of the original travel distance, yielding an effective throughput increase of $21.2/78.8 \approx 26.9\%$. Applied to the 13,000-store network, this translates to approximately 3,500 additional stores serviceable by the same infrastructure.

- Workforce optimization: The same order volume could be fulfilled by approximately 2,049 selectors ($2{,}600 \times 0.788$), freeing roughly 551 positions for redeployment or absorption through natural attrition.

- Selector throughput increase: Each selector's daily throughput is projected to rise from 14 to approximately 18 pallets per shift ($14/0.788 \approx 17.8$, a $1.27\times$ improvement).

12.6 Conclusions

This project developed a three-stage optimization framework for A101's warehouse product placement. The Golden Zone Model identifies high-impact products via frequency and affinity analysis, the Slotting Model assigns products to locations while enforcing pallet stability, and the post-processing phase aligns the layout with store shelf configurations through controlled swaps bounded by a 50-cell budget jointly determined with the industrial advisor. Performance testing showed up to 61.24% improvement in picking distance, with a conservative benchmark of 21.20% formally agreed upon with the industrial advisor for operational projections. The feasibility analysis confirmed deployment readiness, and the pilot at the Esenboğa warehouse was scheduled for July 2026. Future work includes extending the framework to cold-storage and frozen zones, integrating real-time demand forecasting, and measuring field performance after deployment.

Bibliography

A101 (2025). Hakkımızda. https://www.a101.com.tr/hakkimizda/. Erişim tarihi: 21 Ekim 2025.

Chan, F. T. S. and H. K. Chan (2011). Improving store replenishment performance by extracting knowledge from association rules. *Expert Systems with Applications 38*(5), 5894–5900.

Duque-Jaramillo, J. C., J. M. Cogollo-Flórez, C. G. Gómez-Marín, and A. A. Correa-Espinal (2024). Warehouse management optimization using a sorting-based slotting approach. *Journal of Industrial Engineering and Management 17*(1), 133–150.

Gallo, G., P. L. Hammer, and B. Simeone (1980). Quadratic knapsack problems. *Mathematical Programming 12*(1), 132–149.

Van Gils, T., K. Ramaekers, A. Caris, and R. B. M. de Koster (2018). Designing efficient order picking systems by combining planning problems: State-of-the-art classification and review. *European Journal of Operational Research 267*(1), 1–15.

Viveros, P., K. González, R. Mena, F. Kristjanpoller, and J. Robledo (2021). Slotting optimization model for a warehouse with divisible first-level accommodation locations. *Applied Sciences 11*(3), 936.

Appendices

12.A Golden Zone Model Formulation

Symbol	Description
I	set of candidate products
K	set of product categories
S_k	subset of candidate products in category k
P_i	normalized picking frequency of product i
A_{ij}	affinity score between products i and j
C	Golden Zone capacity
Max_k	upper bound on selected products from category k
α, β	weights of frequency and affinity terms
x_i	1 if product i is selected for the Golden Zone; 0 otherwise
z_{ij}	1 if products i and j are jointly selected; 0 otherwise

$$
\begin{aligned}
\max \quad & Z = \alpha \sum_{i \in I} P_i x_i + \beta \sum_{i \in I} \sum_{j > i} A_{ij} z_{ij} \\
\text{s.t.} \quad & \sum_{i \in I} x_i = C \\
& \sum_{i \in S_k} x_i \leq Max_k \quad \forall\, k \in K
\end{aligned}
$$

$$z_{ij} \leq x_i,\ z_{ij} \leq x_j \quad \forall\, i \in I,\ j > i$$
$$z_{ij} \geq x_i + x_j - 1 \quad \forall\, i \in I,\ j > i$$
$$x_i \in \{0, 1\},\ z_{ij} \in \{0, 1\} \quad \forall\, i \in I,\ j > i$$

12.B Slotting Model Formulation

Symbol	Description
I	set of products to be assigned
J	set of eligible storage locations
K	set of durability classes
f_i	picking frequency of product i
D_{jX}, D_{jY}	travel distances from the entry and exit points to location j
G_{ik}	1 if product i belongs to class k; 0 otherwise
Z_{jk}	1 if location j is reserved for class k; 0 otherwise
x_{ij}	1 if product i is assigned to location j; 0 otherwise

$$\begin{aligned}
\min\ Z = \sum_{i \in I} \sum_{j \in J} \big[f_i \cdot (D_{jX} + D_{jY})\big] \cdot x_{ij} \\
\text{s.t.} \quad \sum_{j \in J} x_{ij} = 1 \quad \forall\, i \in I \\
\sum_{i \in I} x_{ij} = 1 \quad \forall\, j \in J \\
x_{ij} \leq \sum_{k \in K} G_{ik} Z_{jk} \quad \forall\, i \in I,\ j \in J \\
x_{ij} \in \{0, 1\} \quad \forall\, i \in I,\ j \in J
\end{aligned}$$

13 Müşteri Deneyiminin İyileştirilmesi İçin Ölçüm, Değerlendirme ve Operasyonel Planlama

Tepe Kurumsal Çözümler

Proje Ekibi

Doğa Aynacı, Zeynep Melis Bakan, Yunus Emre Dermencioğlu, Derviş Mehmet Dinç, Yunus Özcan Malkaya, Ahmet Serdar Yakıcı

Şirket Danışmanı

Melih Günaydın

Veri Analitiği Uzman Yardımcısı

Akademik Danışman

Prof. Dr. Nesim Kohen Erkip

Endüstri Mühendisliği Bölümü

Özet

Müşteri deneyimi verilerinin operasyonel süreçlere entegre edilmesini hedefleyen bu projede, Tepe Kurumsal Çözümler için veri odaklı bir karar destek sistemi geliştirilmiştir. Proje kapsamında, müşteri geri bildirimlerini sayısallaştıran makine öğrenmesi tabanlı bir memnuniyet skorlama yaklaşımı ile bu çıktıları denetim planlama süreçlerine entegre eden matematiksel modeller tasarlanmıştır. Geliştirilen sistem; denetçi atama, dönemsel planlama ve reaktif güncelleme mekanizmalarını bütünleşik bir yapı altında birleştirerek iş yükü dağılımında denge sağlamış, planlama süreçlerini daha uygulanabilir ve kontrollü hale getirmiştir. Ayrıca, müşteri memnuniyetsizliğini erken aşamada tespit ederek proaktif müdahale imkânı sunmuş ve denetim kaynaklarının daha verimli kullanılmasına katkı sağlamıştır.

Anahtar Sözcükler: Müşteri Deneyimi, Karar Destek Sistemi, Atama, Denetim Planlama, Makine Öğrenmesi, Reaktif Planlama.

Measurement, Evaluation, and Operational Planning for Improving Customer Experience

Abstract

This project developed a data-driven decision support system for Tepe Kurumsal Çözümler with the aim of integrating customer experience data into operational processes. Within the scope of the project, a machine learning–based satisfaction scoring approach was designed to quantify customer feedback, and mathematical models were developed to incorporate these outputs into audit planning processes. The proposed system integrates auditor assignment, periodic scheduling, and reactive update mechanisms into a unified framework, improving workload balance and enabling more controlled and feasible planning. In addition, it allows early detection of customer dissatisfaction signals, supporting proactive interventions and more efficient utilization of audit resources.

Keywords: Customer Experience(CX), Decision Support System, Auditor Assignment, Routine Scheduling, Machine Learning,Reactive Scheduling.

13.1 Company and Problem Definition

Founded in 1989, Tepe Kurumsal Çözümler is Turkey's leading provider of integrated facilities and corporate services, operating through six specialized subsidiaries — Tepe Güvenlik, Tepe Tesis Yönetimi, Tepe Gurme, Tepe ISG, Tepe Pro, and Tepe One — that collectively deliver private security, facility management, catering, occupational health and safety, and expense management solutions. With nearly 30,000 employees and a nationwide service network serving over 1000 clients across banking, energy, manufacturing, healthcare, and education sectors.(Tepe Kurumsal Çözümler, 2026)

Tepe Kurumsal Çözümler manages its customer relationships through a CX team supported by 45 auditors — 22 internal and 23 external — who conduct project-based operational audits across nationwide sites. Customer feedback is gathered through surveys, field visits, and complaint records, yet remains scattered across emails, Excel files, and disconnected platforms. Audit schedules are experience-driven, and dissatisfaction signals surface only during monthly management meetings.

Annual visit plans are prepared without explicitly accounting for workload fairness, and equitable coverage across the entire customer portfolio cannot be guaranteed. Survey structures are not standardized across service lines — question sets, scale definitions, and scoring methods vary by company and service type — causing similar feedback to receive different ratings and rendering cross-service comparisons unreliable. The audit visiting mandate is triggered only after observable dissatisfaction signals have already emerged, such as low survey scores, serious complaints, or termination notices. Analysis of conducted audits for single-service customers takes approximately 10 days, while consolidating reports for multi-service customers may extend to two to three months, eliminating the possibility of early intervention.

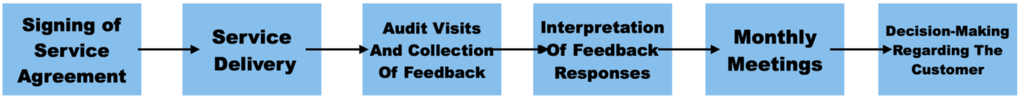

Figure 13.1: Current System Flow Chart.

13.2 Model and Proposed System

The proposed DSS follows the operational sequence used in implementation: *Initial Assignment (MIP)* → *Assignment Heuristic (large instances)* → *Daily Bundle Heuristic* → *Routine Scheduling (year-start / mid-year)* → *Reactive Monthly Update*. All modules run behind the desktop interface; Excel is used as the structured data layer for inputs and outputs.

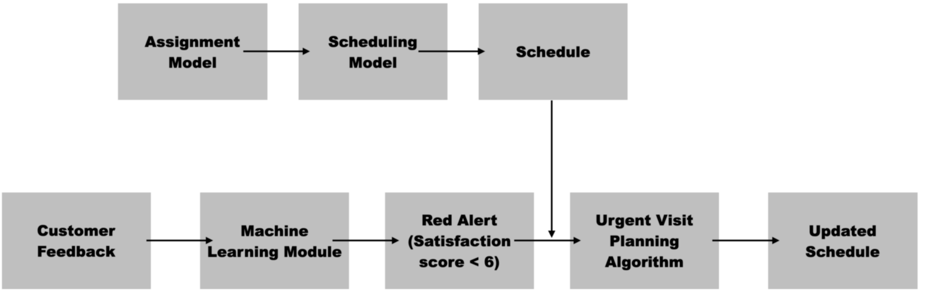

Figure 13.2: System Workflow Diagram.

Machine Learning (Satisfaction Scoring)

Open-ended customer comments are processed via text cleaning and normalization and then scored using a BERT-based satisfaction scoring model. These scores are mapped to a 0–10 scale and combined with numerical survey responses. Customer-level averages and category-level dissatisfaction ratios are then computed. Customers with scores below the threshold are flagged as *Red Alert* and passed to the reactive layer.

Phase I: Auditor Assignment Model (Job-Based Formulation with Workload Fairness)

The Phase I model determines which auditor should be assigned to each predefined audit job. A job may represent a single-customer task or a predefined bundle covering multiple customers. The assignment is made at the job level so that all activities implied by a job are handled by a single auditor. To ensure scalability in large instances, the implementation applies: (i) candidate filtering to reduce A_j, and (ii) time-budgeted solving, where the best feasible incumbent within the allowed time is accepted. The detailed mathematical formulation can be found in Appendix 13.A.

Large-Scale Assignment Heuristic

In cases where the number of projects and auditors becomes very large, the solution time of the assignment model increases significantly. To address this issue, a two-stage heuristic approach has been developed. In the first stage, instead of using all auditors, a smaller core set of auditors is selected and the assignment problem is solved on this subset (seed solution). This stage includes the constraints of the classical assignment model; each project is assigned exactly to one auditor, $\sum_a x_{a,j} = 1$, and auditor workload capacity and operational constraints are satisfied. In this way, the problem size is reduced and a fast and feasible initial solution is obtained.

In the second stage, the projects assigned to each auditor in the core solution are redistributed among all auditors within the group associated with the corresponding auditor. This redistribution is carried out using a heuristic algorithm in which projects are considered in decreasing order of their cost/workload magnitude with respect to the reference auditor, and each project is assigned to candidate auditors in the group according to the following multi-criteria score function:

$$\text{Score}(a, j) = \alpha \frac{L_a + c_{a,j}}{\bar{L}} + \beta \frac{N_a + 1}{\bar{N}} + \theta \frac{c_{a,j}}{\min_{b \in G(j)} c_{b,j}}$$

Here, L_a denotes the current workload of auditor a, N_a denotes the number of assigned jobs, $c_{a,j}$ denotes the project–auditor cost, and $G(j)$ represents the auditor group associated with project j. After the initial assignment, the solution is further improved using a relocation-based local search that minimizes the following objective function to enhance workload and job-count balance:

$$\text{Obj} = \lambda \frac{L_{\max} - L_{\min}}{\bar{L}} + (1 - \lambda) \frac{N_{\max} - N_{\min}}{\bar{N}}$$

This structure is inspired by list-based heuristic approaches used in parallel machine scheduling problems, enabling a balanced and implementable assignment while significantly reducing the solution time (Su et al., 2017).

Daily Bundle Heuristic

In cases where the number of projects is high, manually creating bundles one by one and entering them into the system can lead to significant time loss in the assignment model. To address this issue, an automatic bundle generation heuristic operating under a daily working time constraint has been developed. The objective is both to reduce the manual workload of bundle preparation and to ensure that the total operation time for each bundle does not exceed the predefined daily limit H:

$$\sum_{(i,j)\in R(B)} t_{i,j} + \sum_{j\in B} v_j \leq H$$

Here, $t_{i,j}$ represents all travel times along the route (including travel from the office to the first project, transitions between projects, and the return to the office after the last project), v_j denotes the service time at project j, B represents the set of projects within the bundle, and $R(B)$ denotes the set of route connections associated with the bundle.

In the first stage, projects are sorted based on their polar angles relative to their associated office (home) locations, and initial bundle candidates are generated using a sweep (angular clustering) approach. In the second stage, for each candidate group, the route is initialized with a seed project and the remaining projects are sequentially inserted using a cheapest insertion heuristic. When inserting a project j between two consecutive nodes i and k on the route, the marginal cost is computed as follows:

$$\Delta(i,j,k) = \alpha(t_{i,j} + t_{j,k} - t_{i,k}) + \beta v_j$$

The insertion with the minimum cost is selected, and only those insertions that do not violate the daily time constraint are accepted. After constructing an initial solution, the route sequence is further improved using a 2-opt based local search, minimizing the total travel time:

$$\min \sum_{(i,j)\in R(B)} t_{i,j}$$

Feasible and operational daily bundles that can be directly used in the assignment model are generated automatically, eliminating the need for users to manually construct bundles one by one (Solomon, 1987).

Phase II: Job-Based Audit Scheduling Model with Selective Coverage

This phase plans customer visits over a finite annual planning horizon while satisfying contractual minimum visit requirements and balancing workload across months. The model is executed independently for each auditor after the assignment phase.

The planning unit is a *job*. A job may cover a selective subset of customers. When a job is executed in a given month, only the customers covered by that job are considered visited. This reflects the company's operational practice, where a city visit does not necessarily imply visiting all customers in that city.

The model ensures minimum annual visit fulfillment, supports selective job coverage, balances monthly workload, and discourages repeated consecutive-month visits without making them infeasible when contractual requirements are tight.

The scheduling module supports both year-start planning and mid-year replanning. In mid-year mode, completed visits are read from the existing schedule, deducted from requirements, and only the remaining visits are replanned. The module produces both detailed schedule tables and scheduling KPI sheets, including monthly workload distributions. The detailed mathematical formulation of the scheduling model can be found in Appendix 13.B.

Phase III: Reactive Tasks (Monthly Reactive Update Algorithm)

In Phase III, a monthly reactive planning algorithm is applied to update the annual plan with minimal operational disruption. A reactive visit is an unplanned, urgent audit request that arises outside the routine annual schedule, typically triggered by a low satisfaction score or a customer complaint. At the beginning of each month, the existing plan is updated based on incoming reactive visit requests. As inputs, the current schedule, the list of reactive requests, and system parameters (StartMonth, WeeklyHours, FrozenMonths) are used. Reactive requests are processed in the order they are received (i.e., by submission timestamp), without any priority-based reordering, and are assigned to the currently responsible auditor of the corresponding project; reassignment is not allowed.

For each auditor a and month m, the monthly capacity is computed as

$$\text{MonthlyCap}(a, m) = \text{WeeklyHours}(a) \cdot \frac{\text{DaysInMonth}(m)}{7}.$$

The reactive insertion window is defined as:

$$\text{WindowCap}(a, m) = 0.5 \cdot \text{MonthlyCap}(a, m),$$

which approximately represents a two-week capacity. If the duration of a reactive task exceeds this window, it cannot be inserted into that month.

The algorithm proceeds in three main steps. First, completed tasks are locked (Status = Done), and tasks scheduled before StartMonth but not yet completed are carried over as backlog to the starting month. If capacity is exceeded, these tasks are postponed to future months. Frozen months — months that are locked for planning purposes because operational commitments have already been made — cannot be used for either insertion or postponement.

Second, for each auditor–month pair, the remaining capacity is calculated as

$$\text{RemainingMonthlyCap} = \text{MonthlyCap} - \text{LockedWorkload},$$
$$\text{RemainingWindowCap} = 0.5 \cdot \text{MonthlyCap}.$$

Finally, reactive requests are processed sequentially. For each request, the initial target is StartMonth; if the month is frozen, the next feasible month is considered. If RemainingWindowCap is sufficient, the request is inserted and capacities are updated. If not, cancellable contract tasks (CancelledDueToReactive) are removed within the same month to create capacity. If capacity is still insufficient, the request is either shifted to the next feasible month or postponed further. This process continues until the request is successfully inserted or becomes infeasible.

The planning horizon is effectively unbounded; tasks can be shifted beyond month 12 (e.g., to months 13, 14, etc.), while frozen constraints apply only to months 1–12. After a reactive insertion, planned visits for the same project in the insertion month and the following month are removed and marked as RemovedAsDuplicateAfterReactive.

The algorithm is deterministic (order-preserving), auditor-bound (owner-preserving), sensitive to contract priorities, respectful of frozen months, and non-destructive (updating statuses instead of deleting tasks).

13.3 Validation

All models proposed within the audit planning and assignment system were verified and validated; full technical details are provided in the Final Report. The proposed system was validated across both the machine learning component and the audit planning modules, with a focus on alignment with realistic operational expectations. For the machine learning model, validation was conducted on historical audit records split into training and test sets (see Final Report for exact dates and volumes). Training and test errors were observed to be close, indicating good generalization, and most predictions remain within a narrow error range of actual values.

For the audit planning modules, validation was conducted through multiple scenario runs on both synthetic and company-provided datasets. The assignment model was validated against the company's previous-year audit data under a single-auditor scenario; the resulting weekly workload of approximately 56 hours was operationally plausible given the unplanned nature of prior operations. The bundle heuristic was validated on a real dataset of 218 company projects, where bundling reduced job count by 48.2% and total operation time by 28.0%. The scheduling model's outputs were independently recalculated and matched the model results exactly. The most significant improvement was observed in workload balancing: compared to a random assignment baseline, workload standard deviation decreased from 20.42 to 2.78, the workload gap decreased from 52.87 to 6.98 hours, minimum workload increased from 21.55 to 40.28 hours, and maximum workload decreased from 74.41 to 47.26 hours. When integrated with the scheduling module, workload balance improvements were preserved throughout the planning pipeline. The reactive planning module was validated using company-based scenarios derived from real project data, where all 20 reactive requests in the stress test were successfully placed; the system responded consistently to operational constraints such as frozen periods and changing run months, producing outputs that are both logically consistent and managerially interpretable.

Overall, the system produces consistent and meaningful outputs; however, due to limited geographical differentiation and the absence of historical reactive demand data, the results should be interpreted as a proof of concept.

13.4 Benchmarking and Benefits

The proposed system was benchmarked to evaluate its performance under different operational scenarios.

For the assignment model, benchmarking under bundled and non-bundled settings confirmed model correctness and improved workload balance. Bundling preserved feasibility and produced balanced assignments even in single-auditor scenarios, while demonstrating strong scalability in larger instances (e.g., 45 auditors). Although the total workload reduction was limited (approximately 20 hours), this is mainly due to operational constraints such as the 8-hour limit and short travel distances.

For the planning model, the bundled approach resulted in more balanced monthly workloads and improved schedule feasibility. Compared to the previous non-bundled plan (approximately 2780 hours), the total audit duration was reduced to around 2004 hours, indicating a significant efficiency gain.

For the machine learning module and the reactive planning algorithm, direct benchmarking is not available due to the absence of standardized baseline systems, and results should be considered preliminary due to limited data and the lack of full operational deployment.

Overall, the system improves workload balance and operational efficiency, reduces total audit time, and enables scalable management of increasing audit volumes, providing a strong foundation for data-driven decision making.

13.5 Implementation and Integration

The developed system was evaluated through a controlled pilot study under realistic conditions. Due to data security considerations, the system was deployed in a standalone environment and tested locally, with active involvement of the industrial adviser.

The user interface serves as the central control layer of the system, enabling users to manage data inputs, execute models, and analyze outputs in an integrated and user-friendly environment (see Figure 13.3). The data input module allows users to upload required datasets such as customer locations, job definitions, and operational parameters. The model execution panel enables users to run the assignment, scheduling, and reactive planning modules in a sequential and integrated manner.

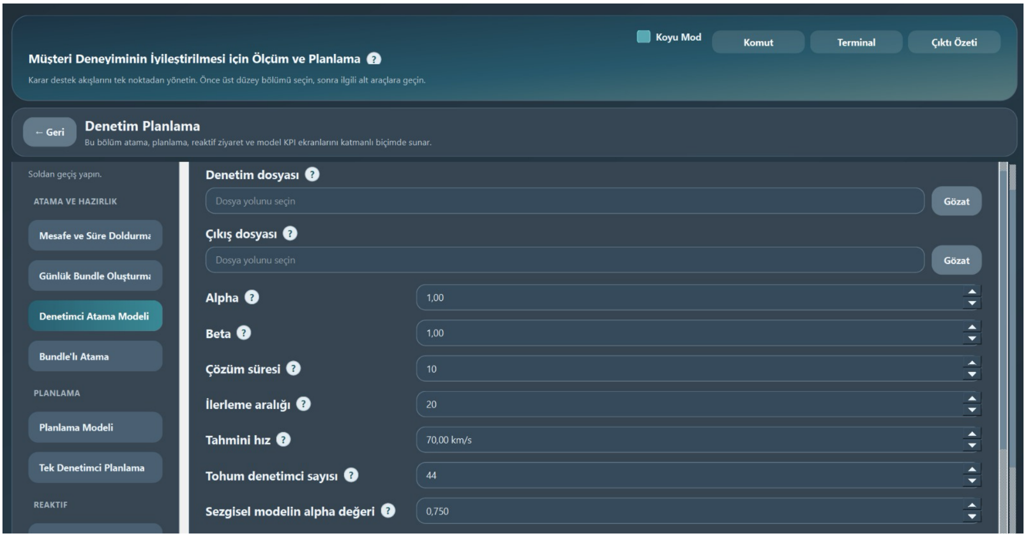

Figure 13.3: Audit planning and parameter configuration interface.

The system operates as a closed-loop structure: assignment outputs are transferred to the scheduling model, and scheduling results, combined with machine learning–based satisfaction scores, trigger the reactive algorithm when necessary. The dashboards provide visual and numerical summaries of key performance indicators, and the system generates structured analytical reports integrating all module outputs for joint evaluation.

The pilot study evaluated system performance under real operational conditions, analyzing the consistency of machine learning–based satisfaction scores with field observations, the responsiveness of the reactive algorithm, and the impact of assignment and scheduling models on workload balance. Throughout the pilot phase, user feedback was collected and improvements were made to the interface design, model parameters, and reporting structure.

13.6 Conclusion

The developed decision support system combines machine learning–based satisfaction analysis with optimization models for assignment, scheduling, and reactive planning, enabling a proactive and data-driven approach to customer experience and operational planning. The pilot study, conducted through the developed user interface, demonstrated that the system can operate effectively under realistic conditions. The industrial advisor confirmed that the system addresses gaps that had been managed through unstructured processes, and expressed that the improvements in workload balance and reactive planning represent a meaningful step toward operationally sustainable audit management.

Bibliography

Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research 35*(2), 254–265.

Su, H., M. Pinedo, and G. Wan (2017). Parallel machine scheduling with eligibility constraints: A composite dispatching rule to minimize total weighted tardiness. *Naval Research Logistics 64*(3), 249–267.

Tepe Kurumsal Çözümler (2026). Tepe kurumsal Çözümler. https://tepekurumsal.com.tr/. Erişim tarihi: 10 Nisan 2026.

Appendices

13.A Assignment Model

Sets and Indices

- A: set of auditors, indexed by a.

- J: set of predefined jobs, indexed by j.
- $A_j \subseteq A$: candidate auditors eligible for job j after candidate filtering.

Parameters

- $\alpha, \beta \geq 0$: objective weights for efficiency and fairness.
- $C_{a,j} \geq 0$: annual workload cost if job j is assigned to auditor a.
- $m_a \geq 0$: optional minimum number of jobs assigned to auditor a.

In the implemented system, $C_{a,j}$ aggregates annual travel effort and annual on-site effort. For single-customer jobs, office-to-site travel is computed. For bundle jobs, the cost is computed along the route legs.

Decision Variables

$x_{a,j} \in \{0,1\}$	1 if job j is assigned to auditor a
$W_a \geq 0$	total annual workload of auditor a
$W_{\max} \geq 0$	maximum workload across auditors
$W_{\min} \geq 0$	minimum workload across auditors

Objective Function

The objective balances overall workload (efficiency) and workload dispersion (fairness):

$$\min \ \alpha \cdot \frac{1}{|A|} \sum_{a \in A} W_a \ + \ \beta \cdot (W_{\max} - W_{\min})$$

Constraints

Exactly-one assignment (with candidate sets).

$$\sum_{a \in A_j} x_{a,j} = 1 \qquad \forall j \in J$$

Workload definition.

$$W_a = \sum_{j \in J:\, a \in A_j} C_{a,j} x_{a,j} \qquad \forall a \in A$$

Max–Min linking (fairness).

$$W_{\max} \geq W_a \qquad \forall a \in A$$

$$W_{\min} \leq W_a \qquad \forall a \in A$$

Optional symmetry-breaking for co-located auditors. If multiple auditors share the same office/city location, they can be ordered to reduce symmetric solutions:

$$W_{a_i} \geq W_{a_{i+1}} \qquad i = 1, \ldots, k-1$$

13.B Scheduling Model

Sets

- C: set of assigned customers.
- J: set of predefined jobs.
- $M = \{1, \dots, 12\}$: set of months in the planning horizon.

Parameters

- R_c^{min}: minimum annual visit requirement for customer c.
- $a_{cj} \in \{0, 1\}$: equals 1 if job j covers customer c.
- k: number of past months with fixed historical decisions.
- $H_{c,m} \in \{0, 1\}$: historical customer visit decision for customer c in month m, defined for $m \leq k$.
- $dist_j \geq 0$: travel time associated with job j.
- $dur_j \geq 0$: audit duration associated with job j.
- $\alpha, \beta \geq 0$: weighting parameters for travel and audit duration.
- $cost_j = \alpha\, dist_j + \beta\, dur_j$: overall workload cost of job j.
- $\lambda \geq 0$: weight associated with workload deviation.
- $p_{gap} \geq 0$: penalty coefficient for consecutive monthly visits.

Decision Variables

$y_{j,m} \in \{0, 1\}$	1 if job j is executed in month m
$x_{c,m} \in \{0, 1\}$	1 if customer c is visited in month m
$W_m \geq 0$	total workload in month m
$W_{\max}, W_{\min} \geq 0$	maximum and minimum monthly workload
$Dev \geq 0$	workload deviation
$u_{c,m} \geq 0$	slack variable for consecutive-visit penalties

Objective Function

$$\min\ \lambda\, Dev\ +\ p_{gap} \sum_{c \in C} \sum_{m=1}^{11} u_{c,m}$$

Constraints

Minimum annual visit requirement.

$$\sum_{m \in M} x_{c,m} \geq R_c^{min} \qquad \forall c \in C$$

Visit definition via jobs (selective coverage).

$$x_{c,m} \leq \sum_{j \in J} a_{cj} y_{j,m} \qquad \forall c \in C, \ \forall m \in M$$

$$\sum_{j \in J} a_{cj} y_{j,m} \leq |J| \, x_{c,m} \qquad \forall c \in C, \ \forall m \in M$$

Fixing historical visit decisions (rolling horizon).

$$x_{c,m} = H_{c,m} \qquad \forall c \in C, \ \forall m \in \{1, \ldots, k\}$$

Monthly workload definition.

$$W_m = \sum_{j \in J} cost_j \, y_{j,m} \qquad \forall m \in M$$

Workload deviation definition.

$$W_{\max} \geq W_m \qquad \forall m \in M$$

$$W_{\min} \leq W_m \qquad \forall m \in M$$

$$Dev = W_{\max} - W_{\min}$$

Soft penalty for consecutive visits.

$$u_{c,m} \geq x_{c,m} + x_{c,m+1} - 1 \qquad \forall c \in C, \ \forall m \in \{1, \ldots, 11\}$$

Müşteri Gelir Tahmini | 14

Hayat Finans

Proje Ekibi
Nazlıcan Akman, Orkun Arslan, Selvi Ecem Civelek,
Arda Noyan Karaşoğlu, Şima Ünver, Göze Yılmaz

Şirket Danışmanı
Alperen Doğan
Kredi Analizi Uzmanı

Akademik Danışman
Prof. Dr. Savaş Dayanık
Endüstri Mühendisliği Bölümü

Özet

Bankacılıkta kritik öneme sahip olan gelir tahmini, doğrudan kredi değerliliğini ve ürün verimliliğini etkilemektedir. Bu proje, geleneksel beyan ve e-Devlet verilerinin ötesine geçerek ek gelir kalemlerini ve ikamet lokasyonunu içeren daha kapsamlı bir model sunmaktadır. Mevcut sistemdeki yanlış tahminler; müşteri memnuniyetsizliğine, artan dolandırıcılık risklerine ve ödeme gücü olan müşterilerin kaybına neden olmaktadır. Geliştirilen model, gelirle ilişkili müşteri özelliklerini ve ikametgah etkilerini korelasyon ve artık kalıntı analizleri ile inceleyerek yeni bir öznitelik seti oluşturmuştur. Geliştirilmiş öznitelik seti makine öğrenmesi yöntemleri ile gelir tahmininde kullanmıştır. Uygulanan yöntemle, gelir tahminlerindeki ortalama mutlak yüzdelik hata (MAPE) %171,2'den %27,6'ya düşürülmüştür. Bu iyileştirme, banka karlılığını artırmış, riskleri minimize etmiş ve müşteri deneyiminde e-Devlet entegrasyonu gibi uzun vadeli verimlilik sağlamıştır.

Anahtar Sözcükler: E-Devlet entegrasyonu, makine öğrenimi, gelir tahmini, dolandırıcılık riskleri, müşteri özellikleri ve korelasyonu.

Customer Income Prediction

Abstract

Accurate income estimation is of critical importance in the banking sector, directly impacting creditworthiness and product valuation. This project proposes a comprehensive model that extends beyond traditional declarations and e-Government data by incorporating additional income streams and residential location dynamics. Inaccurate estimations lead to customer dissatisfaction, heightened fraud risks, and loss of potential revenue from solvent customers. The developed model analyzes the correlation between income-related customer attributes and residential effects through residual analysis to construct an enhanced feature set. Utilizing machine learning techniques, this approach reduced the Mean Absolute Percentage Error (MAPE) from 171.2% to 27.6%. Consequently, the model optimized the bank revenue, mitigated operational risks, and provided long-term improvements in customer experience through seamless digital integration.

Keywords: E-Government integration, machine learning, income prediction, creditworthiness, risk, fraud attacks, customer features.

14.1 About the company

Hayat Finans is the first digital bank in Turkey, offering financial services to both individual and corporate consumers through mobile applications and internet banking since 2022. The institution does not currently provide personal loans or credit cards, but instead offers purchase financing through authorized dealers. Since beginning its operations in 2023, Hayat Finans has acquired more than 400,000 individual customers and over 3,200 corporate clients within just one year, becoming Turkey's largest digital bank in terms of total assets (Hayat Finans Katılım Bankası A.Ş., 2025).

14.2 System Analysis

The company determines the credit limit based on the customer's income, which is automatically retrieved from the system through the mobile application's integration with the E-Government platform. In addition to income verification, each applicant undergoes a comprehensive examination based on their credit information retrieved from the Credit Registry Bureau (KKB), which includes their credit history, existing obligations, and payment performance. The current income model is vulnerable to misleading income information due to the absence of rent or unverified income on the E-Government platform, despite the implementation of digital security measures.

14.2.1 Problem Definition

The current income verification system of Hayat Finans relies predominantly on E-Government integration, which retrieves only SGK 4A, 4B and 4C employment records showing formal salary information. Although this provides a guaranteed income prediction data, many potential customers are lost at this process of the customer application due to the longevity and lack of reliance on this system. This results in lower customer retention and satisfaction rates decreasing the company's revenue due to the loss of opportunities. Furthermore, according to research conducted by Hayat Finans, the residence location of a customer provides a significant indicator of their earnings especially in economically volatile environments like Türkiye. The current income prediction does not integrate the residence location data of the customers which is a potential predictor that is foreseen to significantly affect the income prediction accuracy.

14.2.2 Proposed Solution System

The core idea behind the proposed solution system is to construct a machine learning-based prediction pipeline that learns patterns from historical customer data along with implemented rent by province data as an additional predictor and produces an estimated net monthly income for each applicant, better reflecting real-life financial conditions, including the complex non-linear deviations of income predictors.

Critical Assumptions

The proposed system is built on several key assumptions. First, a customer's credit capacity is fundamentally constrained by their verified income, which aligns with BDDK regulations. Second, financial behavior and credit-related indicators such as payment behavior, account activity, and credit usage contain implicit information about a customer's earning capacity that can be learned from historical data. Third, location-based factors play a significant role in explaining income differences. In an economy like Turkiye, regional economic conditions and cost-of-living differences are reflected in housing and rental values, making a customer's residence location a meaningful proxy for income level. Finally, historical data patterns are assumed to be sufficiently stable to generalize to new customers, and the general relationships between features and income are assumed to hold across the customer base.

Major Constraints

The system is developed under several operational and regulatory constraints. Compliance with the Personal Data Protection Law (KVKK) requires that only verified customer data and anonymized external datasets are used, and no sensitive or unauthorized third-party data sources are incorporated. BDDK guidelines further require that any income estimation model produce results consistent with supporting documentation and remain amenable to auditing, which necessitates maintaining an interpretable model alongside any black-box alternative. Additionally, the target variable itself, formally declared income from E-Devlet, may not fully reflect true earnings, introducing noise into the modeling process and requiring careful preprocessing. The system must also be computationally efficient and scalable for a digital banking environment with a large and growing customer base, meaning model complexity must be balanced with practical implementation considerations.

Objectives

The primary objective of the proposed system is to provide a more accurate and reliable estimation of customer income compared to the current approach, capturing hidden income components and reducing dependence on a single external data source. The natural expected result for Hayat Finans is a more reliable decision making process for the team and an easier acquisition pipeline for the customer.

Solution Approach

The solution is implemented as a structured machine learning pipeline consisting of data preparation, feature engineering, model development, and output generation stages. Data Preparation. The dataset provided by Hayat Finans gathered as of 2025, combines data from E-Devlet, the Credit Registry Bureau (KKB), and Hayat Finans registration records. During preprocessing, income values below the national minimum viable wage of 23,000 TL were excluded as erroneous entries, and values above the E-Devlet reporting cap of 200,000 TL were removed as statistical outliers. Missing values encoded as system placeholders were re-mapped as proper missing indicators statistically handled during modelling. Features with direct leakage from the E-Devlet income field were identified and removed before any modeling step. The dataset was split into training (80%) and test (20%) subsets, with 10-fold cross-validation applied throughout to ensure robust out-of-sample performance estimates.

Feature Engineering. Exploratory data analysis revealed that the income distribution is heavily right-skewed, motivating a logarithmic transformation of the target variable, $Y_i = \log(\text{Income}_i)$, validated through a Box-Cox test. The target distribution after removing the outliers and scaling is shown in Figure 14.1.

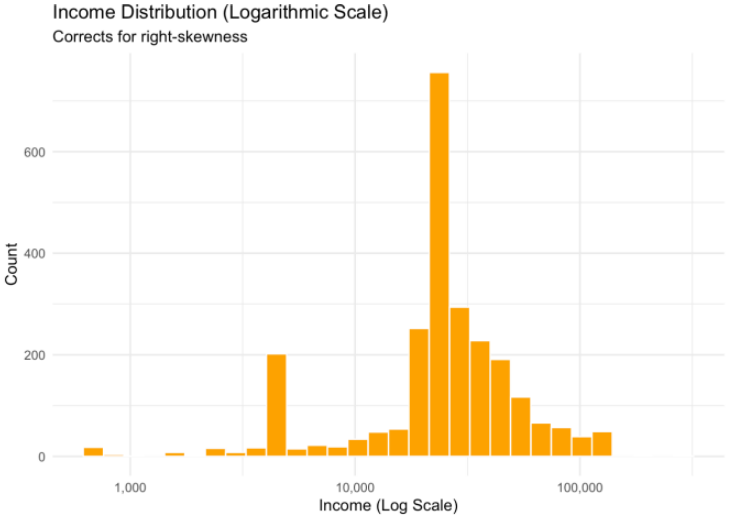

Figure 14.1: Income Distribution on Logarithmic Scale

Scatter plots and LOESS curves showed that several predictors exhibit complex, non-linear relationships with income. Customer age does not increase monotonically with income; earnings generally peak around the mid-30s to early 40s and decline toward retirement. The income distribution through age was shown in Figure 14.2.

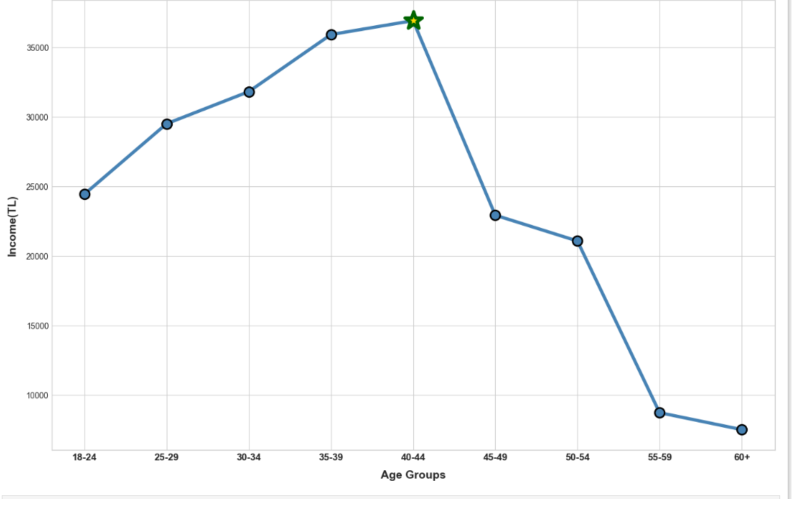

Figure 14.2: Linear Plot of Income Distribution Through Age

This life-cycle pattern was captured by incorporating both a linear and a quadratic age term, as well as an interaction between age and education level to model the experience premium that varies across educational backgrounds. Multicollinearity across the high-dimensional feature set was addressed through Variance Inflation Factor (VIF) analysis and LASSO-guided cluster reduction, consolidating redundant credit limit and delinquency metrics into engineered composite features.

A distinguishing element of the feature engineering process is the integration of location-based economic indicators. Average district and city-level rental prices were retrieved from external platforms including Endeksa, TUIK, and Emlak360, and a positive correlation between provincial rent levels and E-Devlet net income was confirmed empirically. The rent variable was log-transformed and merged with the customer dataset by province, then segmented into income-based quartiles. This engineered feature acted as an effective proxy for regional earning capacity and consistently emerged as one of the strongest predictors of income across all model families evaluated.

Modeling Framework. The modeling workflow was structured along two parallel tracks to balance interpretability with predictive accuracy. In the explainable modeling track, an initial OLS regression served as a diagnostic baseline, revealing severe multicollinearity across the high-dimensional predictor set. LASSO regularization with cross-validated penalty selection was then applied to perform data-driven feature selection, shrinking redundant predictors toward zero and retaining a compact set of stable, informative variables driven by KKB and Hayat Finans operational data. The selected predictors were re-estimated in a reduced OLS model, and residual diagnostics guided further variable transformations. The final explainable model is specified as

$$\begin{aligned} Y_{\log\text{-}income} = \beta_0 + \beta_1 X_{pay} + \beta_2 X_{\log\text{-}rent} + \beta_3(\mathrm{Age})^2 \\ + \beta_4(\mathrm{Age} \times \mathrm{Edu}) + \sum_j \beta_j Z_{ij} + \varepsilon, \end{aligned}$$

where X_{pay} is the log-transformed payment capacity, $X_{\text{log-rent}}$ is the provincial rent proxy, the quadratic age term and its interaction with education capture the income life-cycle, and Z_{ij} denotes the remaining credit behavioral and categorical indicators. An improved specification was further developed by embedding spline terms on key financial variables including credit limits, payment burden, and age, and constructing composite risk signals such as utilization-adjusted credit exposure, to capture non-linearities that a purely additive linear structure cannot adequately represent.

In the predictive machine learning track, several algorithms were trained and evaluated in parallel: Decision Tree, Random Forest, k-Nearest Neighbors, XGBoost, and LightGBM. The performance results of the models were compared in Table 14.1.

Table 14.1: Model Comparison: 10-Fold CV Metrics ($R^2 = \text{cor}^2$)

Model	R^2 **(cor^2)**	**RMSE**	**MAE**
LightGBM	0.4973	0.3118	0.2486
XGBoost	0.4760	0.3200	0.2610
Random Forest	0.4515	0.3335	0.2748
LASSO Regression	0.4241	0.3336	0.2757
OLS Regression	0.3977	0.3421	0.2786
Decision Tree	0.3694	0.3642	0.2882
k-NN	0.1601	0.4201	0.3268

LightGBM, which models income as an iterative set of trees, has been found to be the best performing model. This was an expected result as this was initially trained to be the threshold model. The quite minimalistic errors between ensemble, tree-based and linear algorithms show that the feature engineering process was successful in capturing the complexity of data in a simpler framework.

Model Output and Integration. All log-scale income predictions are back-transformed to Turkish Lira using a smearing correction, which eliminates the retransformation bias that arises from Jensen's inequality. The resulting income estimate feeds directly into Hayat Finans' powercurve system setting the upper bound of credit that can be extended under BDDK regulations.

14.2.3 Verification and Validation

This section presents the verification and validation framework for the proposed income prediction system, evaluating its statistical robustness, predictive reliability, and alignment with Hayat Finans' operational requirements. Verification was carried out through a series of statistical and diagnostic tests to ensure structural integrity and compliance with key modeling assumptions. Initial diagnostics of the baseline OLS model revealed multicollinearity among credit-limit and risk-related variables, which was mitigated by applying a conservative variance inflation factor (VIF) threshold of 10 to maintain model stability. Residual analysis indicated deviations from normality, particularly for lower income levels; these were addressed through data transformations and more flexible modeling approaches. Nonconstant variance tests further confirmed heteroskedasticity, showing that residual variance varied across predicted income levels. This finding motivated the exploration of more flexible models, such as LightGBM, for benchmarking purposes, while retaining the linear modeling framework as the primary approach. Linearity was assessed using partial residual plots as shown in Figure 14.3

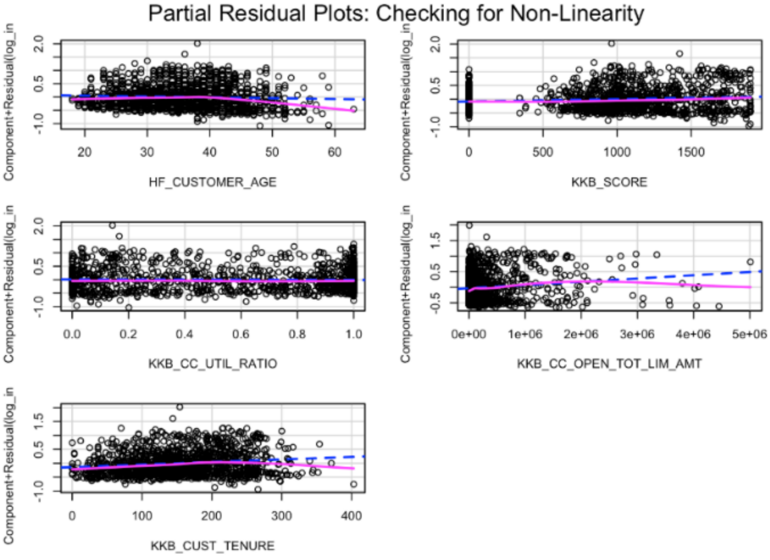

Figure 14.3: Partial Residual Plots: Checking for Non-Linearity

The plots revealed non-linear patterns, including life-cycle effects in age and diminishing returns in credit limits. These were addressed through quadratic and logarithmic transformations. Overall, these verification steps confirm that the model is statistically sound, stable, and theoretically consistent.

Validation of the system focused on the model's ability to generalize to unseen customers while remaining interpretable and operationally feasible within a regulated banking context. Predictive accuracy was evaluated using a stratified 10-fold cross-validation scheme. Although LightGBM achieved the highest performance with the lowest RMSE (0.3118), LASSO regression delivered comparable performance with an RMSE of 0.3336 and was therefore selected as the primary model due to its interpretability and consistency with the linear modeling requirements of the project. Calibration analysis demonstrated that the model successfully preserves income rankings across deciles, with predicted averages closely aligning with observed values. To ensure business interpretability, error metrics were back-transformed into Turkish Lira using a smearing correction to eliminate bias. The results show that the model significantly outperforms a naive mean-based benchmark. Furthermore, segment-specific analysis confirmed that the model maintains consistent performance across different employment statuses, education levels, and residential location. Feature importance was consistent across models as seen in the heatmap of Figure 14.4.

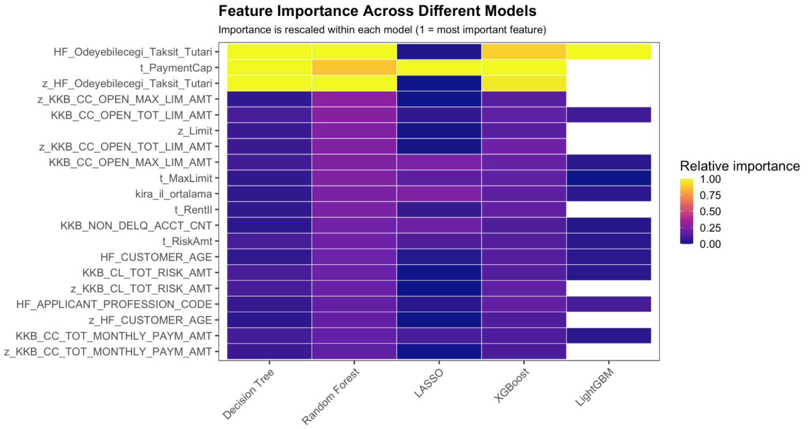

Figure 14.4: Feature Importance Across Different Models

The heatmap shows education level, credit tenure, and installment capacity emerging as key drivers, indicating reliance on economically meaningful factors rather than statistical noise. Overall, the system demonstrates stable and reliable performance under varying input conditions and is suitable for practical implementation within the bank's decision-support framework.

14.2.4 Integration and Implementation

The integration of the proposed income prediction model into Hayat Finans's operational environment was designed around two parallel tracks: alignment with the bank's existing credit decisioning infrastructure, and the development of a user-facing Decision Support System (DSS) to facilitate practical adoption.

On the infrastructure side, Hayat Finans operates its credit evaluation workflow through PowerCurve, a rule-based decisioning platform that determines credit limits based on verified income and risk indicators. The proposed model was designed so that its outputs are directly compatible with this system. Rather than replacing PowerCurve, the model functions as an upstream income estimation layer: it receives the same customer data fields already maintained in the bank's database, including KKB credit history, E-Devlet employment records, and Hayat Finans registration data, and produces a predicted net monthly income that feeds into the existing credit limit calculation. This design ensures that no additional data collection is required from customers, and the model can be deployed without modifying the bank's core decisioning logic. By an initial constraint set by Hayat Finans, the model outputs are positioned as decision support signals rather than automated decisions, meaning credit analysts retain final authority over lending outcomes.

To enable practical use of the model, a Decision Support System prototype was developed as an interactive web application using R Shiny. The DSS was structured with four main components. The components of the DSS were shown in Figure 5.

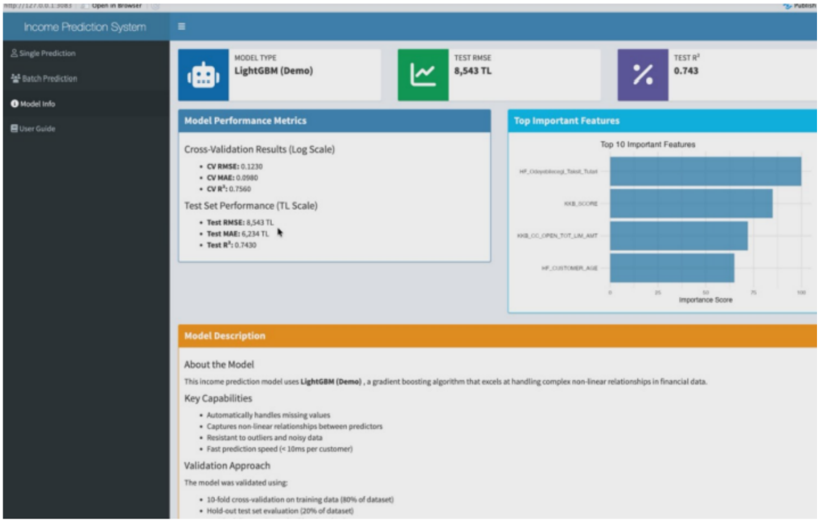

Figure 14.5: DSS Model Info Page

The Single Prediction page allows credit analysts to enter an individual customer's profile through standardized form fields ; covering demographic attributes, credit bureau indicators, and location information and obtain an income prediction with a single action. The interface displays the predicted income alongside key contextual indicators such as the customer's income segment, credit utilization ratio, and payment capacity, enabling analysts to assess predictions within their operational context. For operational scalability, the Batch Prediction page supports processing multiple customers simultaneously by accepting file uploads in CSV or Excel format, running predictions across the entire customer list, and exporting results for internal use. This mode targets scenarios where analysts need to evaluate a portfolio of applicants using the same model configuration. The DSS also includes a Model Information section summarizing the modeling approach, key performance metrics, and the validation setup, serving as a reference for non-technical stakeholders and supporting consistent internal communication about the system's capabilities. A User Guide page provides usage instructions, input requirements, and guidance on interpreting outputs.

The DSS supports three model options: OLS regression, LASSO regression, and LightGBM allowing users to compare predictions across model families. While LightGBM achieved the highest predictive accuracy, LASSO regression was retained as the primary recommended model due to its interpretability and consistency with regulatory requirements for explainable credit decisions. The interface includes a model comparison view that presents side-by-side performance metrics across all three approaches, enabling analysts and management to understand the trade-offs between accuracy and interpretability.

The implementation follows a phased rollout strategy. In the first phase, the model operates in shadow mode alongside the existing system: predictions are generated for incoming customer applications but do not influence actual credit decisions. During this period, E-Devlet income verification continues as before, providing a ground truth benchmark against which model predictions are evaluated on live, unseen customers. This dual-run approach allows the credit analysis team to monitor prediction accuracy, PowerCurve compatibility, and segment-level consistency without introducing operational risk. Qualitative feedback from the credit team and quantitative performance evaluations were collected simultaneously to identify necessary adjustments to input thresholds, feature handling, and output formatting. Based on this feedback cycle, the model parameters and DSS interface was refined before transitioning to operational use in subsequent phases.

14.2.5 Benefits to the Company

The income prediction system provides an income analysis by analyzing existing features with the addition of the location feature based on customer information for better customer service quality, more accurate fraud detection, increased expected revenue from loans and most importantly minimized risk in income prediction. This machine learning approach makes income prediction with a broader feature span that allows flexible decision making as well as analyzing the impact of new features, in our case location. This prediction system decreases the necessity for additional steps in the current system to make sure data on features is reliable while still providing a highly accurate income prediction. The system also decreases the chances of fraud by providing a more detailed customer analysis. Overall, the project allows Hayat Finans to make more accurate and beneficial financial decisions providing a reliable customer base and a tailored customer experience.

14.3 Conclusion

This project successfully met Hayat Finans' goals of reducing reliance on e-Government services in customer revenue forecasting processes and increasing forecast accuracy through data diversity, thanks to the high-performance machine learning models and user-friendly Decision Support System prototype developed. The study, fully aligned with the institution's operational efficiency and risk management expectations, minimized manual processes and enabled the provision of more accurate credit limits. Pilot testing the model with new customer data shows that the reproductive nature of the model allows the reflection of changing socio-economic dynamics in the predictor multiplicator of the linear model.

Bibliography

Hayat Finans Katılım Bankası A.Ş. (2025, March). Hayat finans katılım bankası a.Ş. 2024 yılı faaliyet raporu. Yıllık Faaliyet Raporu. 01.01.2024–31.12.2024 dönemine ait Hayat Finans Katılım Bankası'nın yıllık faaliyet raporudur.

Otomatik Depolama Sistemi Yerleşim ve Çizelgeleme Eniyilemesi | 15

Beko Buzdolabı İşletmesi

Proje Ekibi
Alkım Abanuz, Talha Talip Açıkgöz, Beyza Ceren Gülpınar,
Özge Seydanoğlu, Erdem Ata Tosun, Yusuf Enes Uludağ, Beril Yılmaz

Şirket Danışmanı
Hüseyin Çoker
Kıdemli Uzman Mühendis

Akademik Danışman
Doç. Dr. Firdevs Ulus
Endüstri Mühendisliği Bölümü

Özet

Beko Buzdolabı İşletmesi'nde kullanılan Otomatik Depolama ve Geri Alma Sistemi (AS/RS), operasyonel darboğazlara ve aşırı bekleme sürelerine neden olan tecrübeye dayalı sezgisel yöntemlerle işletilmektedir. Bu problemleri gidermek amacıyla, materyaller için depolama yeri atama ve vinç çizelgeleme süreçlerini eş zamanlı olarak optimize eden bir karar destek sistemi geliştirilmiştir. Çözüm kapsamında, malzeme devir hızlarına dayalı K-ortalamalar kümeleme yöntemi ve Karma Tam Sayılı Doğrusal Programlama (MILP) modelleri kullanılarak gerçek zamanlı karar verme sistemi oluşturulmuştur. Yapılan pilot çalışması sonucunda, sistemin iş çıkarma kapasitesi korunurken ortalama vinç seyahat mesafesinde %25 ve kuyrukta biriken işlerde %41 oranında azalma sağlanmıştır. Geliştirilen yazılım tabanlı çözüm, yıllık yaklaşık 807.600 kWh enerji tasarrufu ve operasyonel verimlilik artışı sağlamaktadır.

Anahtar Sözcükler: Otomatik Depolama ve Geri Alma Sistemi, Depolama Yeri Atama, Vinç Çizelgeleme, Lojistik Optimizasyonu.

Optimization of Storage and Crane Scheduling in Automated Storage System

Abstract

The Automated Storage and Retrieval System (AS/RS) at the Beko Eskişehir Refrigerator Plant operates using experience-driven heuristics which led to operational bottlenecks and excessive waiting times. To address these structural inefficiencies, a solution approach was developed to jointly optimize storage location assignment and crane sequencing. The proposed approach established a real-time decision-making system using turnover-based K-means clustering and Mixed-Integer Linear Programming (MILP) models. Pilot study demonstrates a 25% reduction in average crane travel distance and a 41% decrease in average queue depth. The solution achieves an estimated 807,600 kWh in annual energy savings alongside significant labor efficiency gains.

Keywords: Automated Storage and Retrieval System (AS/RS), Storage Location Assignment, Crane Sequencing, Logistics Optimization.

15.1 Company and System Analysis

15.1.1 Company Description

Beko is Europe's largest home appliances company and a global leader in the industry, operating in more than 55 countries with over 55,000 employees worldwide (Beko Corporate, 2025). Beko combines large-scale manufacturing capability with a strong focus on innovation, efficiency, and sustainability. The Eskişehir Refrigerator Plant, one of Beko's key production facilities, is recognized for its advanced digital infrastructure and Industry 4.0 practices (Home Appliances World, 2021).

15.1.2 Current System Analysis

The Beko Eskişehir Refrigerator Plant operates five production lines supported by an Automated Storage and Retrieval System (AS/RS), which handles approximately 65% of internal material flow. The system consists of 7 aisles with dedicated cranes and stores materials in Eurobox and half-Eurobox containers, which are transferred to production lines through conveyors, monoporters, and elevators. A representation of the AS/RS overview can be seen in Figure 15.1.

Although the AS/RS is a critical component of the production system, it currently operates based on experience-driven heuristics developed for significantly lower production volumes. While originally designed for around 500 units per shift, the system now supports over 4,500 units and handles more than 15,000 material types, and much higher throughput levels, creating a mismatch between system logic and operational requirements.

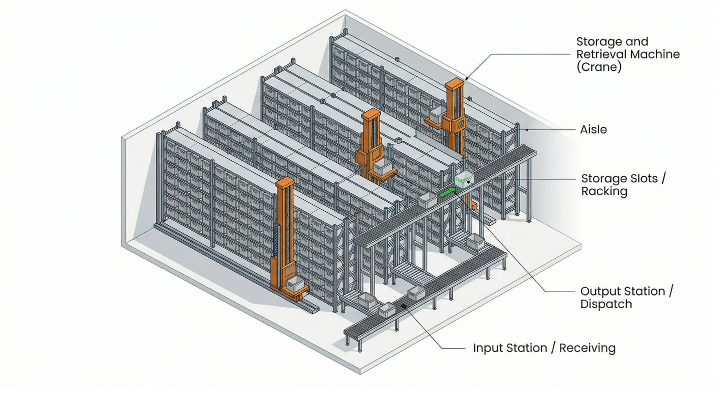

Figure 15.1: AS/RS Overview

15.2 Problem Definition

The AS/RS operates with an automated work-order system where material requests are generated based on workstation needs, and different box combinations can satisfy the same demand. The existing logic assigns cranes based on simple availability, while materials are placed in the first available slot regardless of demand frequency. Furthermore, retrieval tasks are executed chronologically rather than through optimized scheduling. While the workload is distributed across cranes, this process lacks a formal mathematical model and fails to ensure optimal decision-making. The heatmap in Figure 15.2 reveals a nearly uniform distribution of crane visits across the rack; while darker cells indicate higher activity, these high-frequency points are scattered rather than being optimized near the Input/Output (I/O) point to minimize travel distance. In addition, the system experiences fluctuating demand with peak load periods, during which current storage and retrieval policies lead to congestion and delays.

Under the current heuristic-based structure, this leads to inefficiencies in both storage assignment and crane scheduling, especially at high production levels. The proposed approach integrates data analysis, mathematical models, and algorithms to improve system responsiveness, balance crane workload, and reduce unnecessary movements. The scope is limited to optimizing product placement within the AS/RS and crane operations, excluding other material handling systems.

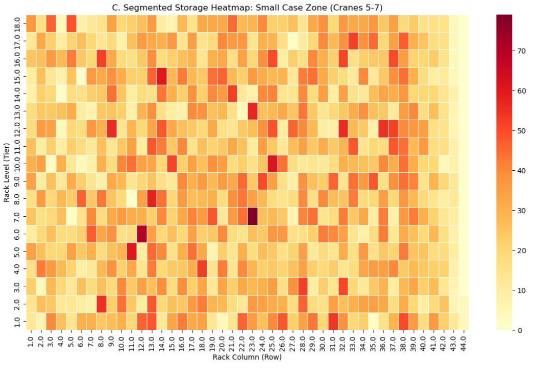

Figure 15.2: AS/RS Crane Visit Frequency Heatmap

15.3 Proposed Solution Strategy

15.3.1 Critical assumptions

The model is built under several simplifying assumptions. Each box is treated as an indivisible unit containing a single material stock code (MSC), and its contents cannot be modified. The AS/RS layout is fixed, including the number and locations of slots and dedicated cranes. Cranes travel at a constant speed, and travel times between slots are calculated using Chebyshev distance, reflecting simultaneous horizontal and vertical movement. Pickup and deposit times are 20 seconds each, and each input point on the single conveyor belt can hold only one box due to the absence of buffer space. Load orders cannot be reassigned between cranes, although they may be re-sequenced within the same queue, whereas unload orders may be reassigned if the item is available in another aisle. All primary material-handling equipment is assumed failure-free, and external supply chain disruptions, production delays, and downtime are excluded from the model.

15.3.2 Major constraints

The AS/RS system handles two box types, namely small and large boxes, whose storage slots are physically separated and therefore cannot be interchanged. Each crane can carry only one box per command cycle and can handle only a single box at a time. Cranes move only parallel to their assigned aisle, without lateral movement, but can serve storage slots on both the left and right sides of that aisle. In addition, each storage slot can hold at most one case at any given time, and every actively managed box must be assigned to exactly one slot within the system. Once a slot is reserved for a load task, it cannot be assigned to another task until the operation is completed or canceled. Finally, an unload task can only be executed if the target slot contains the requested MSC at the time of execution.

15.3.3 Objectives

The overall objective of the proposed solution strategy is to improve AS/RS operational efficiency by reducing work order completion times and balancing workload across aisles. In this way, the system aims to shorten storage and retrieval completion times, improve crane utilization, and increase the number of completed work orders within the existing system structure.

15.4 Solution Approach

15.4.1 Conceptual Model

The proposed solution adopts a layered control system that decomposes the AS/RS scheduling problem into interacting components (Yang et al., 2015). Aisles are modeled as parallel, independent servers, allowing the system to focus on workload balancing and minimizing travel time. The control logic is event-driven, with optimization triggered only by state changes, such as the arrival of new work orders, ensuring both computational efficiency and real-time responsiveness. The overall workflow is illustrated in Figure 15.3.

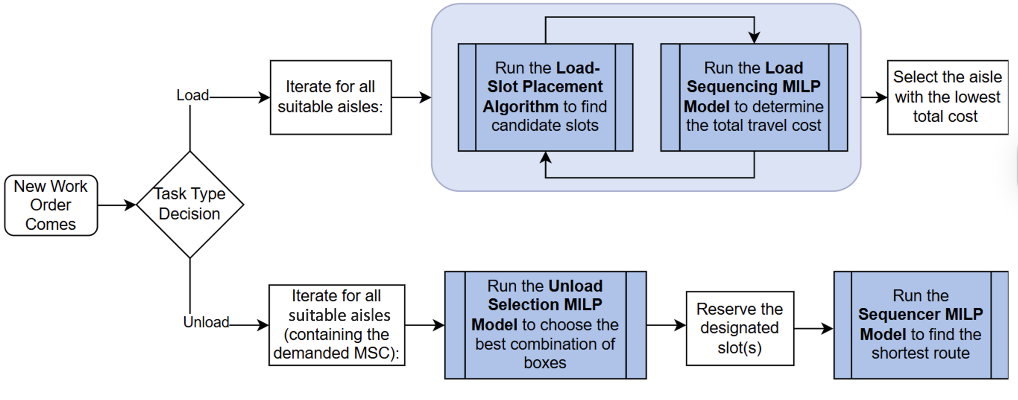

Figure 15.3: Solution Approach Flowchart

When a new work order arrives, the system first determines whether it is a load or unload task. For load operations, feasible aisles are evaluated based on availability and box type compatibility. Candidate storage locations are generated through the Load-Slot Placement Algorithm, and the most suitable aisle, slot, and task sequence are determined using the Load Sequencing Mixed-Integer Linear Programming (MILP) Model. For unload operations, candidate aisles containing the requested MSC are identified, and the Unload Selection MILP Model determines the most appropriate set of boxes to satisfy the demand. If necessary, tasks are divided into smaller sub-orders, selected slots are reserved to prevent conflicts, and the crane's updated task sequence is determined by the Sequencer MILP Model. Detailed descriptions of these models and the algorithm are provided in the following sections.

15.4.2 Mathematical Models

The first mathematical model is the Load Sequencing MILP Model, which minimizes total crane travel distance while determining the best insertion of a new load task into the existing sequence and selecting one candidate slot. The constraints ensure that all fixed tasks are visited exactly once, exactly one candidate slot is chosen, and the resulting tour is feasible and connected. The full formulation is provided in Appendix 15.A.

The Unload Selection MILP Model is formulated as a knapsack-like model that selects the most appropriate subset of boxes containing the requested MSC to satisfy the demand of an unload work order (Polten and Emde, 2022). The objective favors older and fuller boxes while penalizing excessive box retrieval, long travel distances, and surplus quantity. The constraints ensure that the demanded quantity is met and any surplus is explicitly captured. The full formulation is provided in Appendix 15.B.

The Sequencer Model is formulated as a rolling-horizon (MILP) model that determines the execution order of a number H of oldest pending tasks by minimizing total empty travel distance between consecutive tasks. The constraints ensure that each task in the horizon is executed exactly once and that the resulting sequence is feasible and connected. By repeatedly optimizing only the oldest pending tasks, the model preserves First In First Out (FIFO) logic while maintaining computational tractability. The full formulation is provided in Appendix 15.C.

15.4.3 Algorithms

A turnover-based storage assignment methodology was developed to place frequently retrieved items in more accessible locations and thereby reduce travel time. Using the past ten days of historical data, products were grouped according to retrieval frequency through K-means clustering. This approach was preferred over quantile-based classification, since preliminary analysis showed an exponential-like usage distribution that made fixed splits less effective. The methodology is designed to update periodically in order to reflect changing demand patterns.

For load operations, the system uses a heuristic Load-Slot Placement

Algorithm that assigns storage locations based on turnover rate and slot accessibility. For each incoming item, an ideal storage depth is determined from its turnover rate, and all feasible empty slots are scored according to the difference between their actual depth and this ideal position. In this way, high-turnover items are placed closer to I/O points, while lower-turnover items are assigned to less accessible locations.

15.5 Validation

The proposed optimization model was validated against operational data from the company's live AS/RS facility, covering a week from January 2nd to January 9th, 2026.

To ensure a realistic comparison, a 50% normalization factor was applied to real-system timing and queue metrics prior to benchmarking. The value of this factor was determined in consultation with company experts. This adjustment reflects operational noise not captured by the simulation: AS/RS breakdowns and unplanned stoppages, and conveyor congestion at the I/O point. The factor is applied to the real system's baseline so the improvements reported below represent a conservative lower bound.

In terms of performance metrics, the proposed system reduced the average queue depth at the crane conveyors by 52% relative to the adjusted real-system baseline. The improvement in task scheduling further translated into a 40% reduction in average crane waiting time. The item placement logic ensured that frequently requested items were stored closer to the I/O point, cutting average crane travel distance per task by 39%. While queue depth and waiting time metrics are subject to operational variables such as equipment downtime and crane speeds, the 39% reduction in travel distance represents a fundamental structural improvement. This metric provides a deterministic gain in efficiency, independent of external system fluctuations. Furthermore, tasks were distributed evenly across the aisles. Together, these results confirm that the proposed model delivers substantial efficiency gains within the existing hardware infrastructure, even when evaluated under deliberately conservative assumptions.

15.6 Implementation and Pilot Study

The implementation phase was carried out in close collaboration with the company. System logic documentation, annotated pseudo code, technical report, user manual, and detailed flowcharts were prepared and delivered to the company's IT department to support technical integration. Meetings were held to clarify operational constraints, finalize integration requirements, and align the deployment schedule with the factory's existing infrastructure.

Building on the initial one-week validation, a full-scale pilot study was subsequently conducted using one complete month of real operational data, spanning January 2nd to February 4th, 2026. The simulation processed 104,309 work orders across seven crane aisles over a 33-day period. Applying the same 50% conservative adjustment described previously, the proposed model reduced average crane waiting time by 47%, average queue depth by 41%, and average travel distance per task by 25%. These improvements held consistently across all days of the week and all five weeks of the study period, confirming that the results are not driven by any single low-demand period or outlier week.

The main deliverables of the project included optimization modules compatible with the existing AS/RS structure, software code, and a user interface supported by a Python based Decision Support System (DSS) that enabled real-time coordination of storage assignment and crane scheduling within the AS/RS. The DSS was integrated with factory data and background optimization models. The interface was organized into four modules: the System Control Panel for executive KPI tracking and crane monitoring in Figure 15.4; Live Monitoring for corridor heatmaps and slot occupancy in Figure 15.5; Periodical Material Clustering for SKU dataset analysis using K-Means; and Parameter Control for dynamic adjustment of algorithmic weights.

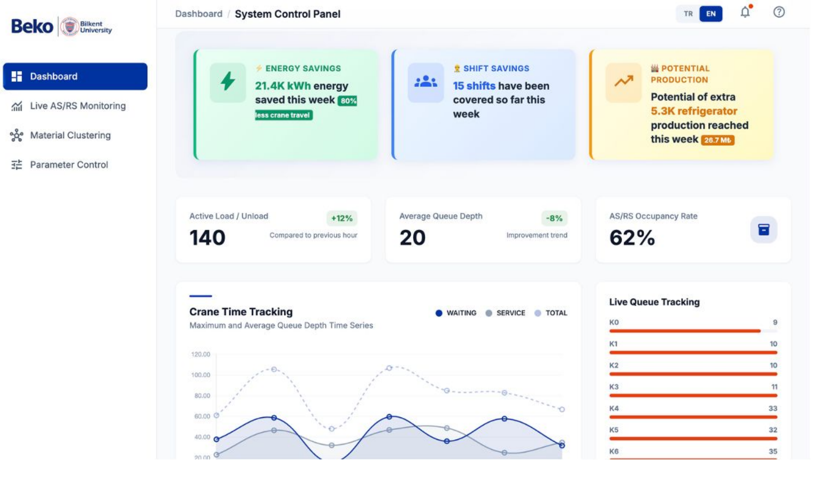

Figure 15.4: Decision Support System: Dashboard

15.7 Benchmarking and Benefits

The proposed framework replaces experience-based heuristics with a data-driven, optimization-based approach for AS/RS operations. Upon receiving

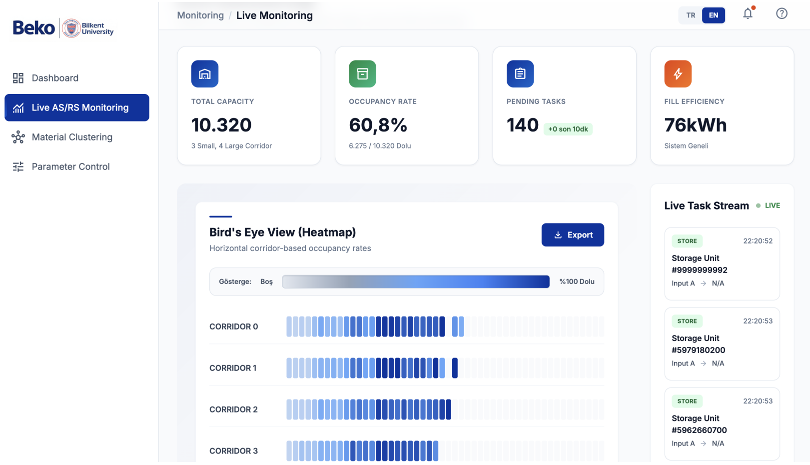

Figure 15.5: Decision Support System: Live Monitoring

a work order, the system activates specific decision modules: storage slots are scored based on demand frequency and proximity to output points, while retrieval tasks are prioritized by balancing stock age, fill levels, and location. By dynamically re-ordering the crane's workload to minimize travel distance and preserve fulfillment sequences, the system significantly reduces unnecessary movement and wait times. Ultimately, this coordinated structure improves operational efficiency by reducing unnecessary crane movement, shortening waiting times, and supporting smoother material flow to production lines.

Benchmark results indicate that the proposed model can significantly improve AS/RS performance compared to the current system. The solution reduces average queue depth by around 41%, average crane travel distance by 25%, and average crane waiting time by 47%, leading to faster response times and more balanced crane utilization. Beyond operational improvements, the project creates a substantial financial impact. Based on annualized estimates, the proposed system provides approximately 807,600 kWh of energy savings, corresponding to nearly $63,810. In comparison, labor efficiency gains are estimated to generate an additional $86,132 in annual value as shown in Table 15.1. Together, these improvements highlight that project delivers both immediate operational benefits and long-term economic value for the company.

15.8 Conclusion

This project has successfully optimized the AS/RS operations at the Beko Eskişehir Refrigerator Plant by integrating storage location assignment and crane sequencing, fully satisfying the organization's expectations for opera-

Table 15.1: Benchmark Results and Economic Impact

Metric	**Current System**	**Proposed Model**	**Improvement**
Average Waiting Time	425 – 476 Sec	∼225 – 252 Sec	47% Reduction
Average Travel Distance	13.6 Unit	10.2 Unit	25% Reduction
Average Queue Depth	5.1 – 6.8 Orders	3.0 – 4.0 Orders	41% Reduction
Financial & Energy Impact	**Annual Savings**	**Economic Value**	**Total Benefit**
Energy Consumption	807,600 kWh	$63,810	
Labor Efficiency	570 Shifts	$86,132	
Total Annual Impact			**$149,942**

tional excellence, as confirmed by positive company feedback. By replacing experience-driven heuristics with a Decision Support System (DSS) utilizing turnover-based K-means clustering and MILP models, the study achieved a 25% reduction in crane travel distance, a 41% reduction in queue depth, and a 47% decrease in crane waiting time. Beyond these technical KPIs, the solution provides approximately $150,000 in annual economic value through energy and labor efficiencies and is currently ready for full-scale integration into the factory's infrastructure. Future work will focus on scaling this model to the other AS/RS unit within the Eskişehir plant and leveraging the identification of breakdown and downtime causes to maximize the system's feeding capacity under ideal operating scenarios.

Bibliography

Beko Corporate (2025). Overview. https://www.bekocorporate.com/company/about-us/overview/. [Online; Accessed October 27, 2025].

Home Appliances World (2021). Arçelik's refrigerator plant included in the wef global lighthouse network. https://www.homeappliancesworld.com/2021/10/13/arceliks-refrigerator-plant-included-in-the-wef-global-lighthouse-network/. [Online; Accessed October 27, 2025].

Polten, L. and S. Emde (2022). Multi-shuttle crane scheduling in automated storage and retrieval systems. European Journal of Operational Research, 302, 892–908.

Yang, P., L. Miao, Z. Xue, and L. Qin (2015). An integrated optimization of location assignment and storage/retrieval scheduling in multi-shuttle automated storage/retrieval systems. Journal of Intelligent Manufacturing, 26(6), 1145–1159.

Appendix: Mathematical Model Formulations

15.A Load Sequencing MILP Model

Notation	Definition
$N = \{0, 1, \dots, m, m+1, \dots, n\}$	Set of all tasks, 0 being the crane's initial position.
$C = \{m+1, \dots, n\}$	Cluster of candidate slots.
$V_{\text{fixed}} = \{0, 1, \dots, m\}$	Set of fixed (existing) tasks.
d_{ij}	Chebyshev distance, travel cost from the end state of task i to the start state of task j.
$X_{ij} \in \{0, 1\}$	Binary variable equal to 1 if the crane moves directly from task i to task j, and 0 otherwise.
$u_i \geq 0$	Continuous auxiliary variable.

$$\text{Minimize} \quad \sum_{i \in N} \sum_{\substack{j \in N \\ j \neq i}} d_{ij} X_{ij} \tag{1}$$

$$\text{s.t.} \quad \sum_{\substack{j \in N \\ j \neq i}} X_{ji} = 1 \qquad \forall i \in V_{\text{fixed}} \tag{2}$$

$$\sum_{\substack{j \in N \\ j \neq i}} X_{ij} = 1 \qquad \forall i \in V_{\text{fixed}} \tag{3}$$

$$\sum_{i \in N \setminus C} \sum_{j \in C} X_{ij} = 1 \tag{4}$$

$$\sum_{i \in C} \sum_{j \in N \setminus C} X_{ij} = 1 \tag{5}$$

$$\sum_{\substack{j \in N \\ j \neq k}} X_{jk} = \sum_{\substack{j \in N \\ j \neq k}} X_{kj} \qquad \forall k \in C \tag{6}$$

$$X_{ij} = 0 \qquad \forall i, j \in C \tag{7}$$

$$u_i - u_j + |N+1| X_{ij} \leq |N+1| - 1 \qquad \forall i, j \in N \setminus \{0\},\ i \neq j \tag{8}$$

$$1 \leq u_i \leq |N+1| - 1 \qquad \forall i \in N \setminus \{0\} \tag{9}$$

$$X_{ij} \in \{0,1\}, \quad u_i \geq 0 \qquad \forall i, j \in N \tag{10}$$

- The objective (1) minimizes total crane travel distance.
- Constraints (2) and (3) ensure that each fixed task is visited once.
- Constraints (4),(5), and (6) ensure that one candidate slot is selected.
- Constraint (7) prevents intra-cluster movements.
- Constraints (8) and (9) eliminate subtours and ensure connectivity.
- Constraint (10) defines the domains.

15.B Unload Selection MILP Model

Notation	Definition
$I = \{0, 1, \ldots, n\}$	Set of all available boxes containing the required MSC.
D	Demanded quantity.
Q_i	Quantity in box i.
A_i	Age score of box i.
Z_i	Distance score of box i.
F_i	Fill score of box i.
w_{age}	Weight for age.
w_{box}	Penalty for each selected box.
w_{dist}	Weight for distance.
w_{fill}	Weight for fill level.
w_{surplus}	Penalty for surplus quantity.
$X_i \in \{0, 1\}$	Binary variable, 1 if box i is selected, 0 otherwise.
$S \geq 0$	Continuous variable, surplus quantity.

$$\begin{aligned}
\text{Maximize} \quad & w_{\text{age}} \sum_{i \in I} A_i X_i + w_{\text{fill}} \sum_{i \in I} F_i X_i - w_{\text{box}} \sum_{i \in I} X_i \\
& - w_{\text{dist}} \sum_{i \in I} (1 - Z_i)\, X_i - w_{\text{surplus}}\, S && (1) \\
\text{s.t.} \quad & \sum_{i \in I} Q_i X_i \geq D && (2) \\
& \sum_{i \in I} Q_i X_i - D = S && (3) \\
& S \geq 0, \quad X_i \in \{0, 1\} \quad \forall i \in I && (4)
\end{aligned}$$

- Objective function (1) favors older and fuller boxes while penalizing long travel, multiple retrievals, and surplus.
- Constraint (2) ensures that total retrieved quantity meets demand.
- Constraint (3) defines surplus as the excess over demand.
- Constraint (4) defines the domains.

15.C Sequencer MILP Model

Notation	Definition
$K = \{0, 1, \ldots, H\}$	Set of tasks in the horizon.
$N = \{0\} \cup K$	Set of all tasks, 0 is the current crane position.
C_{ij}	Travel cost from task i to j.
$X_{ij} \in \{0, 1\}$	Binary variable, 1 if task j follows task i.
$u_i \geq 0$	MTZ auxiliary variable.

$$\text{Minimize} \quad \sum_{i \in N} \sum_{\substack{j \in N \\ j \neq i}} C_{ij} X_{ij} \tag{1}$$

s.t.

$$\sum_{\substack{j \in N \\ j \neq i}} X_{ji} = 1 \qquad \forall i \in N \tag{2}$$

$$\sum_{\substack{j \in N \\ j \neq i}} X_{ij} = 1 \qquad \forall i \in N \tag{3}$$

$$u_i - u_j + |N| X_{ij} \leq |N| - 1 \qquad \forall i, j \in N \setminus \{0\},\ i \neq j \tag{4}$$

$$1 \leq u_i \leq |N| - 1 \qquad \forall i \in N \setminus \{0\} \tag{5}$$

$$X_{ij} \in \{0, 1\}, \quad u_i \geq 0 \qquad \forall i, j \in N \tag{6}$$

- The objective (1) minimizes total empty travel.
- Constraints (2) and (3) ensure each task is executed once.
- Constraints (4) and (5) eliminate subtours.
- Constraint (6) defines variable domains.

16 | Sevkiyat Programlarıyla Uyumlu Depo Dağıtım Planlama Karar Destek Sistemi

Türk Traktör

Proje Ekibi

Zelal Aras, Aysel Erkilet, Berrin Güzel,
Defne Karaman, Lâl Karaoğlu, Celal Özel, Kasım Sangu

Şirket Danışmanları

Murat Göçer, Depo ve Dağıtım Yönetimi Ürün Lideri
Elif Kaplan, Satış Sonrası Tedarik Zinciri Grup Lideri

Akademik Danışman

Prof. Dr. Bahar Yetiş Kara
Endüstri Mühendisliği Bölümü

Özet

Bu proje, TürkTraktör'ün depo dağıtım operasyonlarının anlaşmalı lojistik yüklenici firmanın sevkiyat programına uyumlu planlanmasıyla dağıtım merkezindeki bekleme ve elleçlemenin azaltılmasını hedeflemektedir. Sistem, günlük sipariş toplama ve dağıtım kararlarını bölgesel çıkış günlerine göre planlayan iki aşamalı bir tamsayılı optimizasyon modeliyle ele alınmıştır. Yuvarlanan ufuk yaklaşımıyla bütünleştirilen sistem, aylık dağıtım planını güncel verilerle yenilemekte ve günlük iş emirlerine dönüştürmektedir. Bu yaklaşımla siparişler uygun sevk günlerine atanmakta, günlük dağıtım yapılan bölgelerin dağılımı ve iş emri planı eniyilenmektedir. Geliştirilen sistemle sevkiyat programına uyum %41 artarak %100'e ulaşmış, dağıtım merkezindeki bekleme süresi sıfırlanmış, teslimat gün sapması 8,75'ten 5,8'e düşürülmüş ve kapasite kullanımında %5,44 iyileşme sağlanmıştır.

Anahtar Sözcükler: Dağıtım Planlaması, Lojistik Senkronizasyonu, Tamsayılı Programlama, Yuvarlanan Ufuk Planlaması

Decision Support System for Warehouse Distribution Planning Aligned with Shipment Schedules

Abstract

This project aims to reduce waiting and handling at the logistics service provider's distribution center by aligning TürkTraktör's warehouse dispatch operations with the provider's fixed regional shipment schedule. The system is modeled using a two-stageinteger programming model that plans daily order picking and dispatch decisions based on regional departure days. Using a rolling horizon approach, the system updates the monthly dispatch plan with current data and converts it into daily work orders. As a result, orders are assigned to eligible shipment days, while regional distribution and daily work order plans are optimized. The proposed system improves compliance with the shipment schedule by 41%, reaching 100% compliance, eliminates dwell time at the distribution center, reduces shipment deviation day from 8.75 to 5.8, and improves capacity utilization by 5.44%.

Keywords: Dispatch Planning, Logistics Synchronization, Integer Programming, Rolling Horizon Planning.

16.1 Company Description

Founded in 1954, TürkTraktör (TT) is a leading agricultural machinery manufacturer in Turkey and a major provider of after-sales services. With Koç Holding and CNH as its main shareholders, the company operates under the New Holland and Case IH brands and maintains an important presence in both agricultural machinery and construction equipment. TT has two production facilities in Ankara and Erenler, as well as a spare parts warehouse in Akyurt, Ankara (Türk Traktör, 2025). In 2025, TT accounted for 72% of Turkey's tractor production and 60% of tractor exports to more than 130 countries, while remaining market leader for the 19th consecutive year with a 38.4% share (TürkTraktör, 2026).

16.2 System Analysis and Problem

16.2.1 System Analysis

TT's after-sales operations at the Akyurt Spare Parts Warehouse are carried out through a business-to-business model with 143 authorized spare parts dealers nationwide (Türk Traktör, 2025). Operations begin when customer orders are entered through the SAP portal with either customer-specified or system-assigned requested delivery dates. After stock availability and customer balance checks,work orders are created based on workforce allocation. Each work order consists of picking, packing, and dispatching operations. During thedispatch stage, TT collaborates with a logistics service provider (LSP), which operates under a consolidation-based shipment structure. Invoicing is completed when TT hands over the orders to the LSP, at this stage, the amount is deducted from the customer's balance. Then, they are transferred to the LSP's Ankara Distribution Center (ADC), where they are sorted according to the fixed regional shipment schedule and held until the next scheduled departure. They are then sent to regional distribution centers and delivered to dealers at their final destinations. TT's customers are divided into 17 regions in LSP's service plan. Daily, TT dispatches orders from an average of 14.88 regions. This indicates that, as orders are distributed across almost all regions within a single day, orders from the same regions are split across different days, defined asregional fragmentation. Since the LSP consolidates each region at the ADC, this structure may require additional sorting and handling steps. In addition, TT's dispatch days are not always aligned with the LSP's regional departures. Therefore, some orders arrive at the ADC before their scheduled transfer and remain there until the next departure. This waiting period, defined asdwell time, ranges from 2 to 6 days depending on the shipment schedule. Delivery performance is evaluated based on thedeviation between the requested delivery date and the actual dispatch date. The average delivery deviation in the current system is 8.75 days, indicating that some orders were dispatched almost a week after the requested delivery dates.

16.2.2 Problem Definition

TT's current dispatch structure is based on releasing confirmed and available orders as soon as they become eligible, without incorporating the LSP's fixed regional shipment schedule into dispatch decisions. As a result, dispatch timing is not always fully consistent with LSP's regional departures. In 2024, the compliance rate of TT's dispatches with the LSP's service schedule was 71.03%, indicating that nearly one in three orders experienced dwell at the ADC. This dwell time prolongs the delivery process and directly affects customers, since invoicing is completed when orders are handed over to the LSP. Additionally, extra waiting at ADC might lead to increased sorting and handling, as well as a higher risk of lost items (Göçer, 2025). Accordingly, the project aims to develop a decision support tool that aligns dispatch decisions with the LSP's shipment schedule in order to achieve full schedule alignment, reduce dwell time, improve delivery deviation, and lower regional fragmentation.

16.3 Proposed Solution Approach

16.3.1 Critical Assumptions

The model operates using a rolling horizon approach over a monthly planning horizon. The LSP is assumed to follow the predetermined shipment schedule without delay or revision, and that distribution center locations are fixed and known. Sufficient vehicles are available to meet the planned dispatch schedule. All orders are treated as having the same delivery priority, and all ordered items are assumed to be available in inventory. Daily dispatches are limited by TT's operational capacity. Shipments dispatched before thecut-off, 14:00, are treated as same day dispatches, while those dispatched after the cut-off are assigned to the following day. When daily volumes remain within TT's handling capacity, preparation and shipment are assumed to be completed on the same day.

16.3.2 Major Constraints

Each customer order must be shipped as a single unit, and split deliveries are not permitted. Dispatch activities and departures can occur only during TT's working hours. Dispatch on a given day is allowed only when the planned daily volume exceeds TT's minimum dispatch threshold. If daily dispatch volumes exceed TT's same-day preparation capacity, preparation must begin on the previous day. Orders may be dispatched only within the allowed service window around the requested delivery date. An order can be dispatched on a given day only if the LSP serves that customer on that day, ensuring that no dwell occurs. The number of regions assigned per day is limited to 8 regions.

16.3.3 Objectives

The primary objective is to minimize dwell time and handling within the LSP network through more efficient and coordinated daily dispatch operations at TT. Alignment between TT's dispatch schedule and the LSP's service schedule eliminates dwell time and ensures full compliance. As a result, these requirements are treated as fixed conditions of the model. Therefore, the objective is to assign customer orders to dispatch days based on their regions to reduce fragmentation and improve operational flow. The system assesses the operational performance based on the number of regions served per day, the balance of workload across the planning horizon and work orders, and alignment with customers' requested delivery dates.

16.4 Solution Approach

To incorporate the continuous arrival of orders under real operating conditions, arolling horizon framework is implemented (Sahin et al., 2013). Combined with a planning approach that delays the release of eligible orders to a later dispatch period rather than sending all confirmed orders immediately, this framework provides a more coordinated alternative (Li et al., 2025). In line with this approach, the system assigns orders to feasible dispatch days within an allowable deviation around the requested delivery date.

16.4.1 Conceptual Model

The proposed solution consists of three integrated components: the rolling horizon framework, the Monthly Dispatch Scheduling Model, and the Work Order Planning Model. The visual representation of system's flow is shown in Figure 16.1. The rolling horizon framework updates the plan daily to

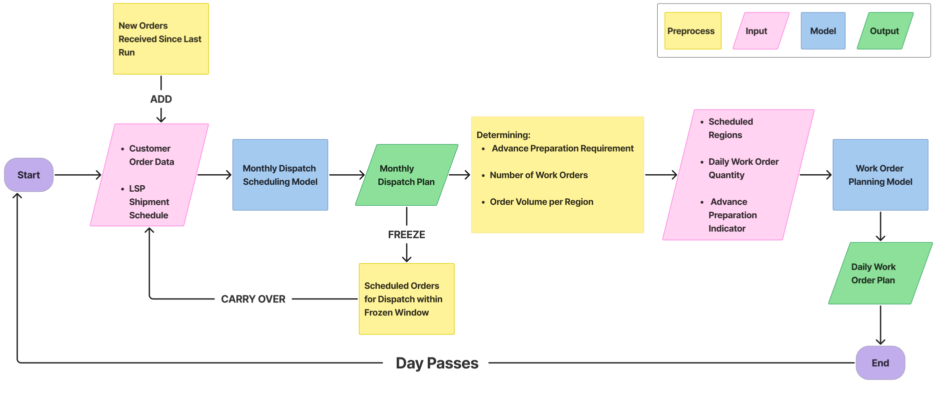

Figure 16.1: The Conceptual Flow Chart of the Proposed System

incorporate new orders while maintaining stability of the near-term dispatch decisions. The planning horizon is divided into frozen, add, and free windows. In thefrozen window, dispatch decisions are fixed. In the add window, existing decisions are preserved, but newly arriving orders can still be incorporated if feasible. In thefree window, the plan for remaining days is re-optimized in each run. Within the rolling horizon structure, the Monthly Dispatch Scheduling Model assigns eligible customer orders to feasible dispatch days over the monthly planning horizon by considering requested delivery dates, TT's operational constraints, and the LSP's regional shipment schedule. The resulting dispatch plan is then transferred to the Work Order Planning Model, which converts near-term dispatch assignments into executable warehouse work orders for the current day and the following day. It also decides whether some of the following day's orders should be prepared after the cut-off on the current day. Together, these three components link long-term dispatch planning with short-term warehouse execution.

16.4.2 Mathematical Model

The proposed solution is formulated as two integer programming models. First, the Monthly Dispatch Scheduling Model determines the dispatch day of each eligible order while minimizing the maximum number of regions served on any dispatch day, as given in objective (16.1); the full formulation is presented in Appendix 16.A. Constraint (16.2) limits the number of regions assigned to the same dispatch day. Constraints (16.3) and (16.4) ensure that each order is assigned to exactly one feasible dispatch day and define its realized dispatch day. Constraints (16.5) and (16.6) capture earliness and lateness relative to the ideal dispatch day. Constraints (16.7) and (16.8) restrict both earliness and lateness by a maximum allowable deviation. Constraint (16.9) calculates the total number of orders dispatched on each day. Constraints (16.10), (16.11), and (16.12) enforce dispatch day activation, daily capacity limits, and minimum order requirements. Constraint (16.13) determines whether a region is served on a given day. Therefore, the model generates a monthly dispatch plan, which serves as an input to the Work Order Planning Model.

Lastly, the Work Order Planning Model minimizes workload differences across work orders, as shown in objective (16.16); the full formulation is presented in Appendix 16.B. Constraint (16.17) assigns each scheduled region to exactly one work order, while constraint (16.18) determines the number of work orders needed according to the planned workload. Constraints (16.19) and (16.20) define the maximum and minimum order limits of work orders. Constraint (16.21) identifies when a part of the following day's dispatch workload must be prepared in advance. Constraint (16.22) ensures that orders prepared in advance are not planned again on their dispatch day. Finally, constraints (16.23) and (16.24) determine how much of the following day's workload will be prepared in advance, subject to the available capacity. In this way, the model creates a balanced daily work plan while also utilizing the after cut-off period for advance preparation when needed.

16.5 Verification and Validation

The proposed system was verified through simplified model analysis, continuity testing, extreme case testing, and rolling horizon stability checks. These analyses confirmed that the models behave as intended under different operating conditions, respond predictably to parameter changes, and preserve frozen dispatch decisions while incorporating new orders beyond the frozen period. Validation was conducted through face validation and operational validation. In face validation, the model structure, key inputs, planning parameters, and generated dispatch and work order plans were reviewed with the Industrial Advisor and found to be consistent with TT's operational practices. Additionally, the solver takes about 17 minutes which was also considered suitable for daily use. Operational validation was performed on order data of May and December 2024 together with the corresponding LSP regional shipment schedules. These months were chosen as May had the highest amount of arrived orders while December had the least amount. Since historical work order data were not available, the comparison was made at the monthly dispatch planning level. The results showed that the proposed system produces a more balanced regional distribution across dispatch days and eliminates dwell time at ADC by aligning dispatch decisions with the LSP's shipment schedule. Therefore, the results indicate that the system produces feasible plans consistent with TT's operations.

16.6 Benchmarking and Benefits

A benchmarking analysis was conducted to compare the proposed decision support system with the current dispatch planning approach under the same operational conditions using 2024 order data. To replicate TT's daily planning process and realistic order arrival times, the system was run daily for the entire year of 2024 within the rolling horizon framework. The comparison demonstrates improvement across all key performance indicators. Under the proposed model, the alignment rate with the LSP's regional shipment schedule increased from 71.03% to 100%, ensuring full shipment compatibility. Accordingly, the dwell time at ADC, which was ranging from 2 to 6 days in TT's system, is completely reduced to zero. Additionally, the average delivery time deviation improved from 8.75 days to 5.8 days, indicating a 33.7% reduction and improved alignment with requested delivery dates. Meanwhile, the average number of regions served per day decreased from 14.88 to 7.90, indicating a more consolidated dispatch structure and reduced potential handling at the ADC. Capacity utilization also improved by 5.44%, reflecting more efficient use of dispatch days and reduced overtime under the proposed planning structure. Taken together, these results show that the proposed model achieves a more coordinated and operationally efficient dispatch system.

16.7 Implementation and Pilot Study

16.7.1 Decision Support System

The proposed decision support system is implemented as an interactive, web-based tool to support TT's dispatch planning process. The system is built usingHTML and CSS, and the models are integrated into it. This structure, as represented in Figure 16.2 allows planners to run the model,

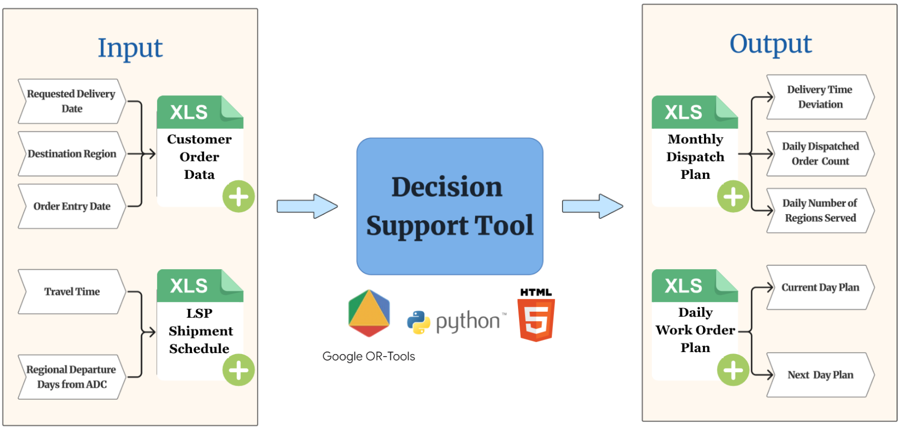

Figure 16.2: Black-Box Representation of Decision Support Tool

review outputs, and evaluate plan performance in a user-friendly environment. The system is designed to operate daily. Its main input is the order data uploaded by the planner. The LSP's fixed regional shipment schedule is embedded in the system, since it does not change frequently (Göçer, 2025). Additionally, the system includes a warm-up mechanism to maintain planning continuity when operations are interrupted due to non-working days. The interface includes a home page that lets the user access either the run page or the KPI dashboard. On the run page, as shown on the right side of Figure 16.3, the user initiates the planning process by selecting the planning start date. The system automatically determines the corresponding horizon end date, and users may optionally include additional operational days, such as overtime days, to reflect temporary planning adjustments. After the order dataset is uploaded, the system executes the optimization models and then generates two outputs in Excel format: a monthly dispatch plan and a work order plan. The KPI dashboard, as shown on the left side of Figure 16.3, summarizes the generated plan through several performance measures. These include average delivery time deviation, average number

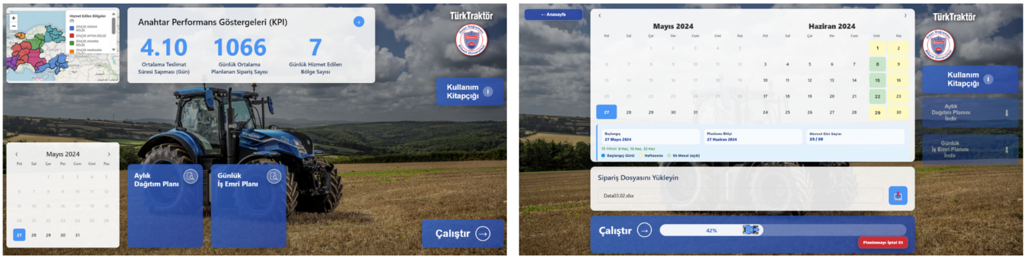

Figure 16.3: Decision Support System: KPI Dashboard (Left) and Planning Panel (Right)

of planned daily orders, and number of regions served per day. The dashboard also includes an interactive map showing the served regions, which can be examined at different zoom levels. In addition, when the delivery time deviation metric is selected, the user can view its distribution through a histogram. Previews of both the monthly dispatch plan and the work order plan are available on this page, allowing users to review outputs directly within the system before downloading them. A user manual is accessible from both the run page and the KPI dashboard to support consistent use of the system and interpretation of the outputs.

16.7.2 Pilot Study

The pilot study, held between April 2-24 2026, by running the system in a company computer in parallel with TT's existing dispatch planning process using live order data. The company was visited on consecutive working days to ensure continuity and integrity of the generated monthly dispatch plans and daily work orders, in line with the rolling horizon framework. The system was not run during the weekend, which was consistent with TT's actual operating practice, and this also showed the warm-up mechanism works as expected. The planning team reviewed the generated dispatch and work order outputs daily alongside the warehouse planning supervisor and logistics coordinator. Planners reported that the system was straightforward to operate and that the user manual adequately supported independent use without requiring technical assistance.

16.8 Conclusion

The project met expectations of TT by aligning operations to the LSP's shipment schedule to reduce dwell time at their distribution center. The proposed system, which combines two integer programming models within a rolling horizon framework, provides optimal monthly dispatch planning and daily work order plans to meet the project's goal. Moreover, the system improved average delivery time deviation, number of regions served per day and capacity utilization to further improve the operations' efficiency. As the stock availability is not tracked in TT's current system, it was not implemented in this project. For further improvements, stock availability could be considered to be part of the system which enables full information flow.

Bibliography

Göçer, M. (2025). Interview and verbal information during warehouse tour at Türk Traktör Akyurt spare parts warehouse. Personal communication. Information provided to the project team during an on-site visit, Türk Traktör, Ankara.

Li, M., X. Qian, M. Li, T. Qu, and Z. He (2025). An order postponement strategy for multi-stage production-logistics synchronization towards zero-warehousing smart manufacturing. *Journal of Manufacturing Systems 82*, 389–404.

Sahin, F., A. Narayanan, and E. P. Robinson (2013). Rolling horizon planning in supply chains: review, implications and directions for future research. *International Journal of Production Research 51*(18), 5413–5436.

Türk Traktör (2025). Biz Kimiz. https://www.turktraktor.com.tr/hakkimizda/biz-kimiz [Accessed: 2025-10-19].

TürkTraktör (2026). Türktraktör 2025 yılı finansal sonuçları açıklandı. Accessed: 2026-04-04.

Appendix: Mathematical Models

16.A Monthly Dispatch Scheduling Model

Table 16.1: Sets and Parameters of Monthly Dispatch Scheduling Model

Notation	Description
K	Set of customer orders
H	Set of potential dispatch days within the planning horizon
R	Set of logistics service provider (LSP) regions
$\Gamma_k \subseteq H$	Set of feasible dispatch days for order k, determined by the LSP schedule
d_k	Requested delivery day of order k
t_k	Transportation time required for order k

$I_k = d_k - t_k$	Ideal dispatch day of order k
β	Maximum allowable deviation from the ideal dispatch day
C	Maximum daily dispatch capacity; for the current implementation it is set to $C = 2500$
L	Minimum number of orders required to activate a dispatch day; for the current implementation it is set to $L = 200$
$a_{kr} \in \{0, 1\}$	Equals 1 if order k belongs to region r, and 0 otherwise
M	Sufficiently large positive constant (Big-M)

Table 16.2: Decision Variables of Monthly Dispatch Scheduling Model

Notation	Description
X_{kh}	1 if order k is dispatched on day h, 0 otherwise
R_{rh}	1 if region r is served on day h, 0 otherwise
W_h	1 if dispatch is active on day h, 0 otherwise
r_k	Dispatch day of order k
$e_k \geq 0$	Earliness of order k relative to its ideal dispatch day
$l_k \geq 0$	Lateness of order k relative to its ideal dispatch day
$N_h \geq 0$	Total number of orders dispatched on day h
$Z \in \mathbb{Z}_+$	Maximum number of regions served on any day

$$
\begin{aligned}
\min \quad & Z && && (16.1)\\
\text{s.t.} \quad & \sum_{r \in R} R_{rh} \leq Z && \forall h \in H && (16.2)\\
& \sum_{h \in \Gamma_k} X_{kh} = 1 && \forall k \in K && (16.3)\\
& r_k = \sum_{h \in \Gamma_k} h X_{kh} && \forall k \in K && (16.4)\\
& l_k \geq r_k - I_k && \forall k \in K && (16.5)\\
& e_k \geq I_k - r_k && \forall k \in K && (16.6)\\
& l_k \leq \beta && \forall k \in K && (16.7)\\
& e_k \leq \beta && \forall k \in K && (16.8)\\
& N_h = \sum_{k \in K : h \in \Gamma_k} X_{kh} && \forall h \in H && (16.9)\\
& X_{kh} \leq W_h && \forall k \in K,\ h \in \Gamma_k && (16.10)
\end{aligned}
$$

$$N_h \leq C \qquad \forall h \in H \tag{16.11}$$

$$N_h \geq L \cdot W_h \qquad \forall h \in H \tag{16.12}$$

$$\sum_{k \in K} a_{kr} X_{kh} \leq M R_{rh} \qquad \forall r \in R,\ h \in H \tag{16.13}$$

$$X_{kh}, R_{rh}, W_h \in \{0, 1\} \tag{16.14}$$

$$r_k, e_k, l_k, N_h \geq 0 \tag{16.15}$$

16.B Two-Day Work Order Planning Model

Table 16.3: Sets and Parameters of Two-Day Work Order Planning Model

Notation	Description
H	Set of dispatch days
R_h	Set of regions on dispatch day h
K_h	Set of potential work orders before cut-off on day h, $K_h \in \{1, 2\}$
c_r^h	Order count of region $r \in R_h$ on dispatch day $h \in H$
α	Maximum preparation capacity for same-day preparation; in the current practice $\alpha = 1750$
k_h	Number of pre-cut-off work orders on day h: $k_h = \begin{cases} 2, & \text{if } \sum_{r \in R_h} c_r^h > 750 \\ 1, & \text{if } \sum_{r \in R_h} c_r^h \leq 750 \end{cases}$
C_{spill}	Maximum preparation capacity for work orders processed after the cut-off time (preparation for the following day's dispatch); in the current practice $C_{\text{spill}} = 750$
SDP_{h+1}	Same-day feasibility indicator: 1 if same-day preparation is feasible for day $h + 1 \in H$, 0 if advanced preparation is required on day $h \in H$
S_{h+1}	$S_{h+1} = \max\left\{0,\ \sum_{r \in R_{h+1}} c_r^{h+1} - \alpha\right\}$: minimum number of orders required to be prepared in advance (excess amount of day $h + 1$)
M	Sufficiently large positive constant (Big-M)

Table 16.4: Decision Variables of Two-Day Work Order Planning Model

Notation	Description
$X_{r,k}^h$	1 if region $r \in R_h$ is assigned to pre-cut-off work order $k \in K_h$ on day $h \in H$, 0 otherwise

p_r^h	1 if region $r \in R_{h+1}$ is selected for advance preparation on day h, 0 otherwise
U_h	Maximum load among all regular work orders planned for day $h \in H$
V_h	Minimum load among all regular work orders planned for day $h \in H$

$$
\begin{aligned}
\min \quad & (U_h - V_h) + (U_{h+1} - V_{h+1}) && && (16.16)\\
\text{s.t.} \quad & \sum_{k \in K_h} X_{r,k}^h = 1 && \forall r \in R_h,\ \forall h \in H && (16.17)\\
& \sum_{r \in R_h} X_{r,2}^h \leq M(k_h - 1) && \forall h \in H && (16.18)\\
& \sum_{r \in R_h} c_r^h X_{r,k}^h \leq U_h && \forall k \in K_h,\ h \in H && (16.19)\\
& \sum_{r \in R_h} c_r^h X_{r,k}^h \geq V_h && \forall k \in K_h,\ h \in H && (16.20)\\
& p_r^h \leq 1 - SDP_{h+1} && \forall r \in R_{h+1},\ h \in H && (16.21)\\
& \sum_{k \in K_{h+1}} X_{r,k}^{h+1} = 1 - p_r^h && \forall r \in R_{h+1},\ h \in H && (16.22)\\
& \sum_{r \in R_{h+1}} c_r^{h+1} p_r^h \geq S_{h+1}(1 - SDP_{h+1}) && \forall h \in H && (16.23)\\
& \sum_{r \in R_{h+1}} c_r^{h+1} p_r^h \leq C_{\text{spill}}(1 - SDP_{h+1}) && \forall h \in H && (16.24)\\
& X_{r,k}^h,\ p_r^h \in \{0, 1\} && && (16.25)\\
& U_h,\ V_h \geq 0 && && (16.26)
\end{aligned}
$$

Sezonsal ve Rassal Talep Altında Dinamik Üretim Planlama | 17

Nesco Gıda

Proje Ekibi
Derin Aktürk, Osmantan Beşikci, Melike Karalar,
Berke Kemahlı, Kenan Arda Köymen, Bahadır Maral, Deniz Yaşar

Şirket Danışmanı
Nazlı Esen Ummansu
Kurucu/CEO

Akademik Danışman
Doç. Dr. Özlem Çavuş İyigün
Endüstri Mühendisliği Bölümü

Özet

Bubble tea alanında üretim yapan ve Nesco Gıda'nın alt markası olan BobaCo, belirsiz ve sezonsal talepler ve üretim kapasitesi kısıtları nedeniyle müşteri talebini etkin biçimde karşılamakta zorlanmaktadır. Bu projenin amacı, şirketin mevcut kapasitesi ile bu talebi daha verimli karşılayabilmesi için bir üretim ve envanter karar destek sisteminin geliştirilmesidir. Ele alınan sistemin hesaplama zorlukları nedeniyle, rassal yaklaşım ile döner çevren mantığını birleştiren bir sezgisel çözüm yöntemi önerilmiştir. Bu yöntem, iki farklı planlama ufkunda çalışan matematiksel modellerin hiyerarşik biçimde bütünleştirilmesiyle sağlanmıştır. Önerilen yaklaşım ve geliştirilen karar destek sistemi, fazla mesai gereksinimi ve birikmiş sipariş ölçülerinde iyileştirme sağlayan yüksek kaliteli ve uygulanabilir üretim planları üretmektedir.

Anahtar Sözcükler: Hiyerarşik Üretim Planlaması, Döner Çevren Yeniden Eniyileme, Rassal Programlama, Uzun Dönem Planlama, Mevsimsel ve Belirsiz Talep

Dynamic Production Planning Under Seasonal and Stochastic Demand

Abstract

BobaCo, a bubble tea producer and sub-brand of Nesco Gıda, faces difficulties in meeting customer demand due to uncertain seasonal patterns and production capacity constraints. This project aims to develop a production and inventory decision support system to enable the company to meet this demand more efficiently under its existing capacity structure. Due to the computational complexity of the system, a heuristic solution approach that combines stochastic modeling with a rolling horizon methodology is proposed. This approach is built on the hierarchical integration of mathematical models operating over two different planning horizons. The proposed approach and decision support system generate production plans that improve overtime and backlog metrics.

Keywords: Hierarchical Production Planning, Rolling Horizon, Stochastic Programming, Long-Term Aggregate Planning, Seasonal and Uncertain Demand, Inventory Build-Ahead, Decision Support System

17.1 Company Description

NESCO is a domestically managed food and beverage company founded in 2016 and headquartered in Ankara, Türkiye, operating with an innovative approach in the industry. From raw material sourcing and product development to production, worldwide sales, and distribution, NESCO manages all major functions of its value chain in-house, allowing it to maintain product quality and respond flexibly to customer needs. It supplies more than 3,500 companies in Turkey and exports to over 30 countries, serving both domestic HORECA and international B2B customers. Its fastest-growing brand, BobaCo, is among the first brands to manufacture bubble tea on an industrial scale in Türkiye, with 25 distinct popping boba flavors (Nesco Gıda Sanayi ve Ticaret A.Ş., 2025). The products are primarily supplied through a B2B model to cafés, coffee chains, restaurants, hotels, and beverage brands, while some formats are also offered in retail packaging.

17.2 System Analysis and Problem

17.2.1 System Analysis

The BobaCo production system consists of four interrelated products: Plastic Tub Boba, Semi-Finished Boba in Tub, Cup Bubble Tea, and Can Bubble Tea. Plastic Tub Boba is sold directly as a finished product, mainly for the domestic market, whereas Semi-Finished Boba in Tub serves only as an input in the production of Cup and Can Bubble Tea. Production is organized on two lines. The first line produces Plastic Tubs and Semi-Finished Boba, which follow the same production flow but differ in formulation, making them operationally similar yet distinct for planning. The second line produces Cup and Can Bubble Tea, but due to pasteurization constraints, these two products cannot be produced on the same day. In addition, once a line is assigned to a product, it cannot switch to another on the same day, so each line can process at most one product per day.

17.2.2 Problem Definition

In previous years, the company faced pronounced demand peaks during the summer season. However, meeting this demand was difficult due to limited production capacity. Although product shelf life allows inventory to be carried over multiple months, limited storage space in the previous facility prevented the company from building inventory ahead. With the transition to a new facility, storage limitations have eased, creating the opportunity to build inventory in advance. This shift makes it essential to develop a planning approach that coordinates production capacity and inventory decisions over time. The key challenge is to balance holding and backlogging costs while maintaining production feasibility for both raw materials and finished goods under seasonal and uncertain demand. The company prepares its production plan and schedule semi-manually by evaluating market insights, past experiences, and previous demand data. Since demand, particularly in the domestic market, is both seasonal and difficult to predict precisely, a reliable and mathematically grounded decision support system is needed to support production planning under uncertainty. The objective of this decision support system is to generate a production plan that determines the appropriate production and stocking levels of different product types, minimizing holding and backlogging costs during the high-demand season while satisfying system requirements. Accordingly, the problem extends beyond simple capacity allocation and requires coordinated production and inventory decisions under seasonal and uncertain demand.

17.3 Proposed Solution Approach

17.3.1 Critical Assumptions

There were several critical assumptions made in order to create the models effectively and efficiently. Semi-finished Boba and Tub Boba are considered as different products. Each production line can produce at most one product type per day. Backlogging is allowed in the system to preserve feasibility;

however, delays within two weeks are treated as non-critical, whereas delays beyond two weeks are considered more severe and are penalized more heavily to reflect the risk of order cancellations.

17.3.2 Major Constraints

Production lines cannot switch between products within the same day. There are two production lines: the first line produces Plastic Tub Boba and Semi-Finished Boba, while the second line produces Cup and Can Bubble Tea. Total production in each period is limited by the available workdays, excluding overtime. Since Semi-Finished Boba is not directly demanded, it is used as an input in the production of Cup and Can products. Semi-Finished Boba may also be stored for future use. The remaining products are directly demanded by end customers and may also be stored. Storage area of the facility is limited, meaning that the total area occupied by stored products cannot exceed the available capacity. In addition, raw material availability is constrained by procurement lead times and minimum order quantities. Backlogging is allowed, but delays exceeding 12 days are treated as critical through an additional penalty structure.

17.3.3 Objectives

The primary objective is to determine a long-term production and inventory strategy that minimizes holding and backlog costs while ensuring adequate inventory coverage for the high-demand season. Since backlog costs are higher than inventory holding costs, the planning approach is designed to discourage excessive delay while avoiding unnecessary inventory buildup.

17.3.4 Conceptual Model

The main model is formulated as a multi-period production and inventory planning structure that represents how products flow through production, inventory, and demand fulfillment over time. Within this structure, production decisions determine period-by-period output, inventory allows available stock to be carried into future periods, and backlog captures unmet demand that is postponed to later periods.

The model also captures the structural dependency among products within the system. While Plastic Tub Boba, Cup Bubble Tea, and Can Bubble Tea are directly demanded products, Semi-Finished Boba is not directly demanded and instead serves as an input for Cup and Can Bubble Tea. Accordingly, the system is represented as an integrated planning structure in which production and inventory decisions for different products must remain coordinated and consistent over the planning horizon.

17.3.5 Heuristic Solution Method

Hierarchical Heuristic Approach

The initial "full-scale" model was formulated at a daily resolution over a one-year horizon and tailored to the facility's specific operational restrictions and characteristics to reflect the production system accurately. However, solving a fully integrated, daily, full-year model within a reasonable run-time becomes computationally challenging due to the size of the decision space. Therefore, a hierarchical heuristic approach is adopted by decomposing the planning problem into two coordinated models: a long-term aggregate production planning model given in Appendix 17.A and a short-term production planning model given in Appendix 17.B, implemented with a rolling-horizon mechanism (Bitran et al., 1982). This hierarchical coordination between the long-term and short-term models is illustrated in Figure 17.1.

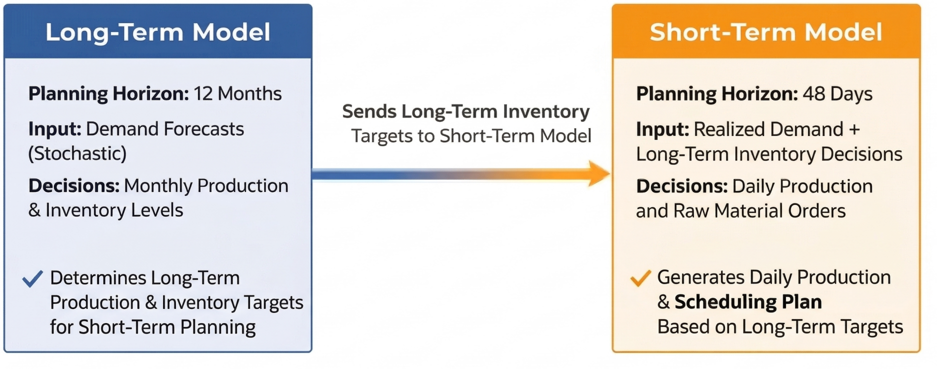

Figure 17.1: Hierarchical Models Diagram

Stochastic Long-Term Model

The long-term model determines strategic inventory targets over a one-year horizon at a monthly level. It identifies in which months production should be increased to prepare for peak-season demand and sets targeted monthly inventory levels. The long-term model optimizes these targets by minimizing expected total cost, including inventory holding and backlog-related penalty components, while maintaining consistency with major system limits and planning rules (A.1).

To account for demand uncertainty, the long-term model is formulated as a two-stage stochastic program, in which demand for the first two months is treated as deterministic, while demand in the remaining periods of the planning horizon is modeled as random (Birge and Louveaux, 2011). Based on the deviation margins provided by the company, an uncertainty band is defined around the company's baseline forecast, and multiple demand scenarios are generated accordingly. Equal probabilities are assigned to each scenario, and the long-term model evaluates the expected total cost across the scenario set. This scenario-based formulation produces monthly inventory targets that are more resilient to uncertainty and reduces the limitations of relying on a single forecast.

Short-Term Production Planning Model

The short-term model uses the inventory decisions for the first two months obtained from the long-term model as input and translates the monthly inventory targets into operationally feasible daily decisions within an eight-week planning horizon (B.11). In particular, it generates a detailed daily production plan that aligns with the facility's line-level constraints and converts the higher-level inventory targets into executable production quantities, inventory updates, and short-term fulfillment decisions (B.2). It also generates raw material order decisions to ensure material availability and maintain consistency between short-term production requirements and procurement needs (B.7),(B.8). In this way, the proposed heuristic approach combines daily production planning in the short-term model with inventory targets generated by the long-term model, resulting in a hierarchical decision support structure that improves peak-season planning while maintaining computational and operational feasibility.

Rolling-Horizon Structure

The two models operate under a rolling-horizon structure: the long-term model is re-run on a monthly basis to refresh monthly inventory targets as new information becomes available, while the short-term model is updated weekly to produce an actionable near-term schedule.

17.4 Validation

Using the company's actual input data and forecast structure, the rolling-horizon decision support system produced planning patterns consistent with the company's operational logic. In particular, during peak-demand periods, the system builds inventory ahead and later uses this inventory to control backlog. It was also observed that when forecasts are updated or available workdays decrease, the rolling-horizon structure responds quickly by revising build-ahead and production decisions, making the system more responsive to changing operating conditions. As shown in Figure 17.2, the model initially built inventory in anticipation of higher forecasted demand, and later reduced production when realized demand turned out to be lower than expected. This behavior is well aligned with the company's planning needs, since the main challenge in the current system is not only meeting seasonal demand, but also reacting in a timely and structured manner to changing production conditions.

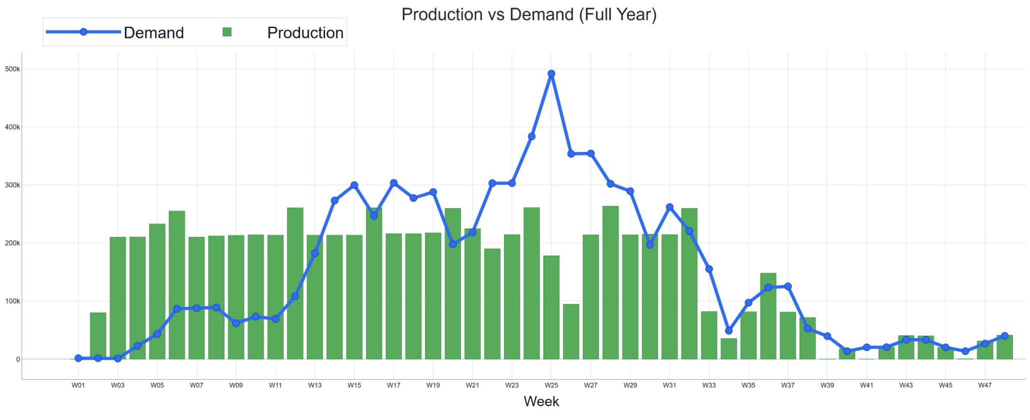

Figure 17.2: Production and Demand Levels Throughout the Year

Both the rolling-horizon structure and the stochastic long-term component contribute meaningfully to the company's ability to prepare for future demand scenarios. While the rolling-horizon structure improves the company's ability to respond to demand changes in a timely and consistent manner, the stochastic long-term model generates more protective inventory targets under uncertainty. In the yearly rolling-horizon implementation, this more cautious inventory policy increased the average holding-cost component by 94.2%, while reducing the average backlog-cost component by 46.4%. This result is consistent with the model structure, since backlog costs are very large and the stochastic model reacts to uncertainty by keeping more inventory in order to reduce the risk of future backlog. Accordingly, although the holding-cost component increases, the stochastic approach reduces the average total cost by 42.0%. These values should be interpreted as model-based performance measures rather than direct accounting costs, but they still indicate that the stochastic policy provides a more robust balance between inventory protection and backlog risk under uncertainty.

17.5 Benchmarking and Benefits

Benchmarking analysis was conducted to compare the proposed decision support system with the company's current planning approach under the same demand inputs. In this comparison, the current system was represented by a 45-day backward allocation rule. Since the company responds to delays exceeding 12 days through overtime usage, the benchmarking focuses on backlog risk and overtime dependency. The following KPIs were used in the comparison: Backlog Under Cancellation Risk, Overtime Days, and Overtime Delivery Rate (the share of total demand that must be recovered through overtime production).

Under the original 2026 forecast, the current system required 11 overtime days and created 380,709 units of backlog under cancellation risk out of a total demand of 6,330,210 units. This corresponds to 6.01% of total demand being under cancellation risk. Under the same forecast, the proposed system eliminated overtime need completely and produced no residual backlog.

To test the robustness of the proposed system under uncertainty, 10 demand scenarios were generated around a monthly forecast. Under these scenarios, the current system required between 17 and 37 overtime days, with an average of 28.2 days, and created 8,082,588 units of backlog under cancellation risk over a total demand of 64,721,203 units, corresponding to a weighted backlog-under-cancellation-risk ratio of 12.49%. In contrast, the proposed system achieved zero backlog and zero overtime in 7 of the 10 scenarios. In the remaining 3 scenarios, it required just 1 overtime day and generated limited residual backlog. Overall, its average overtime requirement was 0.3 days, and its total remaining backlog was 49,330 units, corresponding to 0.08% of total demand. A summary of the comparative performance of the current and proposed systems is presented in Table 17.1.

Table 17.1: Scenario-based results for the current and proposed systems

Metric	**Current System**	**Proposed System**
Average overtime days across 10 scenarios	28.2	0.3
Weighted backlog-under-cancellation-risk ratio	12.49%	0.08%
Scenarios with zero backlog and zero overtime	0/10	7/10

These findings show that the company's current planning logic is highly sensitive to demand fluctuations and frequently requires reactive overtime to recover delayed orders. The proposed rolling-horizon decision support system performs much more robustly by building inventory ahead when necessary and limiting backlog accumulation beyond the 12-day tolerance window. Therefore, the main benefit to the company is not only a reduction in backlog risk, but also a substantial reduction in overtime dependency under both the base forecast and the generated demand scenarios. Numerically, this means that the proposed system reduces the average overtime requirement from 28.2 days to 0.3 days and achieves zero backlog and zero overtime in 7 out of 10 generated demand scenarios.

17.6 Implementation and Pilot Study

The decision support system was integrated into the company's weekly production planning process through a structured four-week pilot study. In coordination with company representatives, an interface was developed in line with the firm's existing reporting practices, which facilitated adoption by planners. Figure 17.3 presents the user interface of the decision support system together with a sample weekly production schedule generated by the system. The first week was conducted under the close supervision of the project team, while the following weeks were carried out directly by the company within its regular planning routine. On a weekly basis, the tool generated recommended production quantities and raw material order decisions based on updated demand information and current operational inputs. These recommendations were reviewed every Monday for feasibility and shop-floor consistency, and the finalized plans and manual adjustments were documented in weekly decision records.

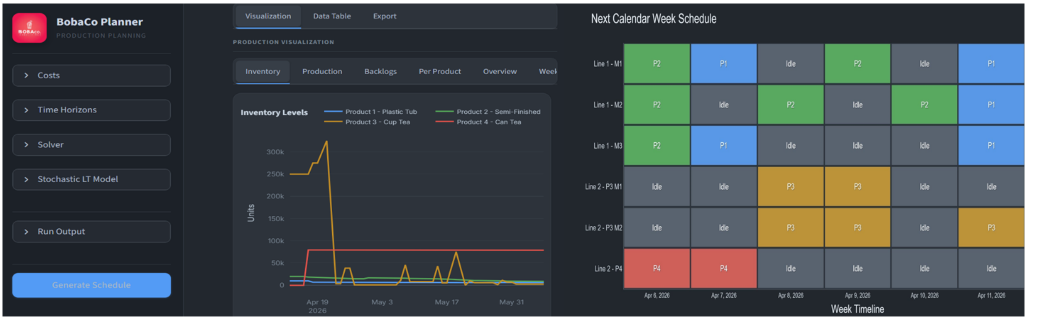

Figure 17.3: System interface and sample weekly schedule output.

The pilot study confirmed that the system could support proactive inventory build-ahead in response to forecasted peak-season demand. In particular, the system led the company to hold inventory that it had not originally planned to keep, but that was necessary to prepare for the high demand anticipated in the 2026 summer forecast. Weekly inventory quantities were recorded to monitor stock levels and their evolution throughout the pilot period. At the end of the study, the recorded decisions and inventory data indicated that the system's outputs were compatible with weekly planning requirements and operational practice. Overall, the pilot met the company's main expectations by demonstrating both operational practicality and the suitability of the system for regular use.

17.7 Conclusion

The project met the company's expectations by developing a decision support system that improves production planning under seasonal and uncertain demand. The proposed structure combines hierarchical production planning, a stochastic long-term model, a short-term production planning model, and a rolling-horizon mechanism to produce plans that are both operationally feasible and responsive to changing conditions. The validation results showed that the system supports proactive inventory build-ahead and reacts quickly to forecast updates and changes in available workdays. In addition, the yearly rolling-horizon implementation results indicated that the stochastic approach provides a more robust balance between inventory protection and backlog risk under uncertainty, reducing the average total cost by 42.0%. This relatively great improvement is mainly due to the model structure, where backlog costs are much higher than holding costs, so carrying additional inventory helps avoid much larger backlog-related losses.

The practical value of the system was also demonstrated through benchmarking and pilot use. Compared with the company's current planning logic, the proposed approach substantially reduced overtime dependency, lowering average overtime requirement from 28.2 days to 0.3 days and achieving zero backlog and zero overtime in 7 out of 10 generated demand scenarios. In addition, the four-week pilot study showed that the system could be incorporated into the company's weekly planning routine and could support inventory positioning for the high demand anticipated in the 2026 summer forecast. Future work may focus on improving the scenario generation structure through observations collected over longer periods, allowing demand uncertainty to be represented in an even more realistic way.

Bibliography

Birge, J. R. and F. Louveaux (2011). *Introduction to Stochastic Programming* (2 ed.). Springer.

Bitran, G. R., E. A. Haas, and A. C. Hax (1982). Hierarchical production planning: A two-stage system. *Operations Research 30*(2), 232–251.

Nesco Gıda Sanayi ve Ticaret A.Ş. (2025). About Us – BobaCo. https://thebobaco.com.tr/pages/hakkimizda/ [Online; accessed October 27, 2025].

Appendices

17.A Stochastic Long-Term Model

Sets:

Symbol	Explanation
K_1	Set of Plastic Tub and Semi-Finished Boba products
K_2	Set of Cup Bubble Tea and Can Bubble Tea products
T	Set of planning periods (months), $T = \{1, \ldots, 12\}$
Ω	Set of scenarios

Parameters:

Symbol	Explanation
$d_{i,t}$	Demand of product $i \in K_1 \cup K_2$ in period $t \in \{1, 2\}$
$d_{i,t,\omega}$	Demand of $i \in K_1 \cup K_2$ in period $t \in \{3, \ldots, 12\}$ under scenario $\omega \in \Omega$
a_t	Number of workdays in period $t \in T$
b_i	Backlog cost for product $i \in K_1 \cup K_2$
c_1	Daily unit capacity of production line 1
$c_{2,i}$	Daily capacity of production line 2 for product $i \in K_2$
r_i	Unit usage area of product $i \in K_1 \cup K_2$
u_{i2}	Number of product 2 used in the product $i \in \{3, 4\}$
h_i	Holding cost of product $i \in K_1 \cup K_2$ for one unit per period

Decision Variables:

Symbol	Explanation
$P_{i,t}$	Production of product $i \in K_1 \cup K_2$ in period $t \in \{1, 2\}$
$I_{i,t}$	Inventory of product $i \in K_1 \cup K_2$ at the end of period $t \in \{1, 2\}$
$B_{i,t}$	Backlog of product $i \in K_1 \cup K_2$ in period $t \in \{1, 2\}$
$P_{i,t,\omega}$	Production of $i \in K_1 \cup K_2$ in period $t \in \{3, \ldots, 12\}$ under scenario $\omega \in \Omega$
$I_{i,t,\omega}$	Inventory of $i \in K_1 \cup K_2$ in period $t \in \{3, \ldots, 12\}$ under scenario $\omega \in \Omega$
$B_{i,t,\omega}$	Backlog of $i \in K_1 \cup K_2$ in period $t \in \{3, \ldots, 12\}$ under scenario $\omega \in \Omega$

Objective Function:

$$\min \sum_{t=1}^{2} \sum_{i \in K_1 \cup K_2} (b_i B_{i,t} + h_i I_{i,t}) + \sum_{\omega \in \Omega} \sum_{t=3}^{12} \sum_{i \in K_1 \cup K_2} p_\omega \left(b_i B_{i,t,\omega} + h_i I_{i,t,\omega}\right) \tag{A.1}$$

Constraints:

Production Capacity Constraints:

$$\sum_{i \in K_1} \frac{P_{i,t}}{c_1} \leq a_t \quad , \quad \sum_{i \in K_2} \frac{P_{i,t}}{c_{2,i}} \leq a_t \qquad \forall t \in \{1, 2\} \tag{A.2}$$

$$\sum_{i\in K_1}\frac{P_{i,t,\omega}}{c_1}\le a_t \quad , \quad \sum_{i\in K_2}\frac{P_{i,t,\omega}}{c_{2,i}}\le a_t \qquad \forall t\in\{3,\ldots,12\},\ \forall\omega\in\Omega \quad \text{(A.3)}$$

Inventory Balance for Semi-Finished Boba:

$$I_{2,t-1}+P_{2,t}-\sum_{i\in K_2}u_{i2}P_{i,t}=I_{2,t} \qquad \forall t\in\{1,2\} \quad \text{(A.4)}$$

$$I_{2,t-1,\omega}+P_{2,t,\omega}-\sum_{i\in K_2}u_{i2}P_{i,t,\omega}=I_{2,t,\omega} \qquad \forall t\in\{3,\ldots,12\},\ \forall\omega\in\Omega \quad \text{(A.5)}$$

Inventory–Backlog Balance for Finished Products:

$$I_{i,t-1}+P_{i,t}-(d_{i,t}+B_{i,t-1})+B_{i,t}=I_{i,t} \qquad \forall t\in\{1,2\},\ \forall i\in K\setminus\{2\} \quad \text{(A.6)}$$

$$I_{i,t-1,\omega}+P_{i,t,\omega}-(d_{i,t,\omega}+B_{i,t-1,\omega})+B_{i,t,\omega}=I_{i,t,\omega} \quad \forall t\in\{3,\ldots,12\},\ \forall\omega\in\Omega,\ \forall i\in K\setminus\{2\} \quad \text{(A.7)}$$

Storage Capacity Constraints:

$$I^{\max}\ge\sum_{i=1}^{4}r_iI_{i,t} \qquad \forall t\in\{1,2\} \quad \text{(A.8)}$$

$$I^{\max}\ge\sum_{i=1}^{4}r_iI_{i,t,\omega} \qquad \forall t\in\{3,\ldots,12\},\ \forall\omega\in\Omega \quad \text{(A.9)}$$

Non-negativity Constraints:

$$P_{i,t},\ I_{i,t},\ B_{i,t},\ P_{i,t,\omega},\ I_{i,t,\omega},\ B_{i,t,\omega}\ge 0 \qquad \forall i\in K,\ \forall t\in T,\ \forall\omega\in\Omega \quad \text{(A.10)}$$

17.B Short-Term Production Planning Model

Sets:

Symbol	**Explanation**
K_1	Set of Plastic Tub and Semi-Finished Boba products
K_2	Set of Cup Bubble Tea and Can Bubble Tea products
T	Set of days in the planning horizon
S	Set of raw materials
M	Set of machines on Line 1, $M=\{1,2,3\}$
M_3	Set of machines for product 3 on Line 2, $M_3=\{1,2\}$

Parameters:

Symbol	Explanation
b_i	Daily unit backlog cost of product $i \in K_1 \cup K_2$
b_i^{12}	Daily unit backlog cost of product $i \in K_1 \cup K_2$ whose due date has exceeded at least 12 days
n_{is}	Number of raw material $s \in S$ used in the product $i \in K_1 \cup K_2$
LT_s	Lead time of raw material $s \in S$
κ_m	Daily capacity of machine $m \in M$ on Lines
u_{2i}	Number of product 2 used in the product $i \in K_2$
$d_{i,t}$	Customer demand of product $i \in K_1 \cup K_2$ at day $t \in T$
$d_{i,1}^D,\ d_{i,2}^D$	Long-term inventory decision parameters
h_i	Unit holding cost of product $i \in K_1 \cup K_2$
h_s^r	Unit holding cost of raw material $s \in S$
v_i	Inventory area usage of one unit of product $i \in K_1 \cup K_2$
$w_{i,t}$	Total demand of product $i \in K_1 \cup K_2$ from day $t-12$ to $t \in T$
moq_s	Minimum order quantity for raw material $s \in S$

Decision Variables:

Symbol	Explanation
$P_{i,t}$	Number of product $i \in K_1 \cup K_2$ produced at day $t \in T$
$I_{i,t}$	Number of product $i \in K_1 \cup K_2$ in inventory at day $t \in T$
$I_{s,t}^r$	Number of raw material $s \in S$ in inventory at day $t \in T$
$R_{s,t}$	Number of raw material $\in Ss$ ordered at day $t \in T$
$Z_{s,t}$	1 if raw material $s \in S$ ordered at day $t \in T$, 0 otherwise
$B_{i,t}$	Backlog of product $i \in K_1 \cup K_2$ at day $t \in T$
$B_{i,t}^{12}$	Backlog of product $i \in K_1 \cup K_2$ at day $t \in T$ whose due date has exceeded at least 12 days
$X_{m,i,t}$	1 if machine $m \in M$ produces $i \in \{1,2\}$ at day $t \in T$, 0 otherwise
$X_{m,3,t}$	1 if machine $m \in M_3$ produces 3 at day $t \in T$, 0 otherwise
$X_{4,t}$	1 if Line 2 produces product 4 at day $t \in T$, 0 otherwise

Objective Function:

$$\min \sum_{t\in T} \sum_{i\in K_1\cup K_2} \left(b_i B_{i,t} + b_i^{12} B_{i,t}^{12} + h_i I_{i,t}\right) + \sum_{t\in T}\sum_{s\in S} h_s^r I_{s,t}^r \tag{B.1}$$

Constraints:

Line Assignment Constraints:

$$X_{m,1,t} + X_{m,2,t} \le 1 \quad , \quad X_{m,3,t} \le 1 - X_{4,t} \qquad \forall m \in M,\ \forall t \in T \tag{B.2}$$

Production Definitions:

$$P_{1,t} = \sum_{m\in M} \kappa_m X_{m,1,t} \quad , \quad P_{2,t} = \sum_{m\in M} \kappa_m X_{m,2,t} \qquad \forall t \in T \tag{B.3}$$

$$P_{3,t} = \sum_{m\in M_3} \kappa_m^3 X_{m,3,t} \quad , \quad P_{4,t} = \kappa^4 X_{4,t} \qquad \forall t \in T \tag{B.4}$$

Inventory Balance Constraints:

$$I_{2,t-1} + P_{2,t-1} - \sum_{i \in K_2} u_{2i} P_{i,t} = I_{2,t} \qquad \forall t \in T \quad \text{(B.5)}$$

$$I_{i,t-1} + P_{i,t} - (d_{i,t} + B_{i,t-1}) + B_{i,t} = I_{i,t} \qquad \forall t \in T,\ \forall i \in (K_1 \cup K_2) \setminus \{2\} \quad \text{(B.6)}$$

Raw Material Balance and Ordering Constraints:

$$I^r_{s,t} = I^r_{s,t-1} + R_{s,(t-LT_s)} - \sum_{i \in K_1 \cup K_2} n_{is} P_{i,t} \qquad \forall s \in S,\ \forall t \in T \quad \text{(B.7)}$$

$$R_{s,t} \geq moq_s \cdot Z_{s,t} \qquad \forall s \in S,\ \forall t \in T \quad \text{(B.8)}$$

Inventory Area Constraints:

$$I^r_{\max} \geq \sum_{s \in S} v_s I^r_{s,t} \quad , \quad I^{\max} \geq \sum_{i \in K_1 \cup K_2} v_i I_{i,t} \qquad \forall t \in T \quad \text{(B.9)}$$

Critical Backlog Constraint:

$$B^{12}_{i,t} \geq B_{i,t} - w_{i,t} \qquad \forall i \in K_1 \cup K_2,\ \forall t \in T \quad \text{(B.10)}$$

Long-Term Target Satisfaction Constraints:

$$I_{i,24} \geq d^D_{i,1} \quad , \quad I_{i,48} \geq d^D_{i,2} \qquad \forall i \in K_1 \cup K_2 \quad \text{(B.11)}$$

Domain Constraints:

$$X_{m,1,t},\ X_{m,2,t},\ X_{m,3,t},\ X_{4,t} \in \{0,1\} \qquad \forall m \in M \cup M_3,\ \forall t \in T \quad \text{(B.12)}$$

$$Z_{s,t} \in \{0,1\} \qquad \forall s \in S,\ \forall t \in T \quad \text{(B.13)}$$

$$P_{i,t},\ I_{i,t},\ B_{i,t},\ B^{12}_{i,t},\ I^r_{s,t},\ R_{s,t} \geq 0 \qquad \forall i \in K_1 \cup K_2,\ \forall t \in T,\ \forall s \in S \quad \text{(B.14)}$$

Sistem Karmaşıklığının Ölçülmesine Dayalı Karar Destek Sistemi | 18

Beko Bulaşık Makinesi İşletmesi

Proje Ekibi
Alp Deniz Çabuk, Barış Uras Durmuş, Tuğrul Kutlu,
Umut Özcan, Deniz Özin, Süeda Eslem Sümer, Emirhan Yavuz

Şirket Danışmanı
Beytullah Uyulur
Karmaşıklık Yönetimi Lider Mühendisi

Akademik Danışman
Prof. Dr. Alper Şen
Endüstri Mühendisliği Bölümü

Özet

Sincan Organize Sanayi Bölgesi'nde yer alan Beko Bulaşık Makinesi Fabrikası, yüksek ürün çeşitliliği, sınırlı parça ortaklığı ve sık değişen mühendislik talepleri nedeniyle yüksek derecede yapısal ve işletimsel karmaşıklıkla mücadele etmektedir. Bu durum, üretim hattında gizli maliyetlere, verimlilik kayıplarına ve parça eksikliğinden kaynaklanan tıkanmalara yol açmaktadır. Bu proje kapsamında, fabrikadaki karmaşıklığı ölçmek, analiz etmek ve en iyilemek amacıyla veri odaklı bir karar destek aracı olan "OptiCost" geliştirilmiştir. Bu karar destek aracı, ürün ağacı verilerini temel almakta ve makine öğrenmesi algoritmaları kullanarak farklı ürün konfigürasyonlarının normalize edilmiş adam-saat gereksinimlerini ve duruş riski seviyelerini tahmin ederek fabrikadaki karmaşıklığın kontrol altına alınmasını hedeflemektedir.

Anahtar Sözcükler: Üretim Karmaşıklığı, Maliyet Azaltımı, Makine Öğrenmesi, Üretim Durma Riski, Karar Destek Sistemi.

A Decision Support System Based on the Measurement of System Complexity

Abstract

The Beko Dishwasher Factory, located in the Sincan Organized Industrial Zone, struggles with a high degree of structural and operational complexity due to high product variety, limited part commonality, and frequent engineering change requests. This situation leads to hidden costs on the production line, efficiency losses, and blockages caused by part shortages. Within the scope of this project, a data-driven decision support tool named "OptiCost" has been developed to quantify, analyze, and optimize the complexity within the factory. Based on bill of materials (BOM) data, this tool utilizes a machine learning algorithm to predict the normalized man-hour requirements and blockage risk levels of different product configurations, ultimately aiming to bring factory complexity under control.

Keywords: Manufacturing Complexity, Cost Optimization, Machine Learning, Blockage Risk, Decision Support System.

18.1 Company Description

Beko is a home appliances company owned by Koç Holding (Beko Corporate, 2026; Koç Holding, 2026). It has its parent and owned companies which also belonged to Koç Holding, such as Arçelik. The company has a workforce of more than 50,000 employees in total, and several factories located in many regions. Some of these regions are Europe, Asia, Africa, and Middle East(Beko Corporate, 2026). While it has a total of 22 parent, owned or limited license brands, Beko became the corporate brand among these brands as a part of its globalization strategy (Bloomberg HT, 2024).

Beko's Ankara Dishwasher Production Facility is one of the 45 production facilities worldwide (Arçelik Global, 2024). There is also a R&D center of the company located in the facility. The facility is located in Sincan, and it has been built on a 109,000 m2 area. 50,580 m2 of this area is used for production, R&D, and warehouse buildings (Google, 2026). In the factory, there are approximately 1200 different dishwasher types that are produced, and the number of components used for the production is approximately 4000.

18.2 Current System and the Problem

18.2.1 System Analysis

Beko's dishwasher production is managed through a structured product-module-component hierarchy. In this system, each product consists of functional modules (e.g., engines, pumps, door assemblies), which are further divided into the smallest material units, or components (e.g., screws, small parts). This hierarchical architecture is subject to constant evolution due to frequent Engineering Change Requests (ECRs) and the continuous introduction of new product configurations.

18.2.2 Problem Definition

The primary challenge in the production system is escalating complexity, driven by the following factors:

- **Growing Variety:** Increasing stock keeping unit (SKU) variety and design volatility make it difficult to track their cumulative impact on operational performance.
- **Lack of quantitative complexity assessment:** There is no systematic,data-driven metric to quantify the complexity added by new designs or engineering changes.
- **Operational Transparency Gaps:** High diversity obscures the direct correlation between product structure and its effect on costs, lead times, and resource allocation.
- **Low Commonality:** Insufficient part-sharing across variants increases purchasing costs and supply chain risks, reducing production stability.

The absence of a quantitative complexity assessment tool leads to measurable inefficiencies in both production and cost-management processes.

18.3 Proposed Solution Strategy

18.3.1 Solution Approach

The proposed framework begins by automating BOM-based structural metrics that were previously calculated manually. Using Bill of Materials (BOM) data, the system computes key indicators such as Unique Part Ratio, Common Part Ratio (CPR), total number of unique components, SKU count, average part count per SKU, and module distributions. These metrics are dynamically updated and form the basis of the analysis. The framework evaluates a given production plan that involves a subset of SKUs using the same workflow. The selected product mix is converted into BOM-based structural metrics at the scenario level. For historical production data, observed blockage outcomes are used, while for a future production plan, blockage risk is estimated using a trained classification model. Normalized man-hour (NMH) values are used as a workload indicator. If available, they are used directly; otherwise, an XGBoost-based model predicts NMH for new or unseen configurations based on product structure. These values are then used in blockage risk evaluation. Blockage is modeled as a classification problem with three levels: LOW, MID, and HIGH.

For every production plan, a system-level complexity score is computed by combining workload and structural factors. A system-level complexity score is computed by combining:

- SHAP-based workload contribution ($\times 0.1$)
- CPR deviation from 0.60
- Scaled Unique Part Ratio ($\times 10$)
- Normalized SKU and component counts
- Average parts per SKU ($\times 0.1$)

The final score is normalized to a 0–100 scale for easy interpretation and comparison. Finally, the system allows direct evaluation of new product configurations. In such cases, NMH is predicted based on structure and can be used independently or integrated into further analysis.

18.3.2 Mathematical Models

Input Representation

Each SKU or production scenario is represented using structural information derived from the Bill of Materials (BOM). At the component level, configurations are encoded as a binary part-existence vector:

$$x^{(i)} = [x_1^{(i)}, x_2^{(i)}, \ldots, x_d^{(i)}]^T, \quad x_j^{(i)} \in \{0, 1\} \tag{18.1}$$

where $x_j^{(i)} = 1$ indicates that component j is present in configuration i, $x_j^{(i)} = 0$ otherwise and d is the total number of distinct components. This representation allows the model to directly capture the structural composition of each product.

In addition, scenario-level features such as SKU count, number of unique components, average parts per SKU, and Common Part Ratio (CPR) are computed from the BOM and used as model inputs .

Normalized Man-Hour Prediction Model

Production complexity is modeled using the Normalized Man-Hour (NMH), which reflects the relative workload required to produce a SKU. NMH values are scaled between 0 and 1 to ensure comparability. An XGBoost regression model is used to estimate NMH:

$$\hat{y}^{(i)} = F(\mathbf{x}^{(i)}) = \sum_{k=1}^{K} f_k(\mathbf{x}^{(i)}), \quad f_k \in \mathcal{F} \tag{18.2}$$

Here, $\mathcal{F}$ is the set of regression trees, and each f_k maps the input vector to an output value. The model is trained by minimizing the following objective function:

$$\mathcal{L} = \sum_{i=1}^{n} \ell\left(y^{(i)}, \hat{y}^{(i)}\right) + \sum_{k=1}^{K} \Omega(f_k) \tag{18.3}$$

where $\ell\left(y^{(i)}, \hat{y}^{(i)}\right)$ measures prediction error and $\Omega(f_k)$ penalizes model complexity to avoid overfitting.

SHAP-Based Explainability for NMH

To improve interpretability, SHAP (Shapley Additive Explanations) is applied to the NMH model. For a given input $\mathbf{x}^{(t)}$, the prediction can be decomposed as:

$$F(\mathbf{x}^{(t)}) = \phi_0 + \sum_{j} \phi_j^{(t)} \tag{18.4}$$

Here, ϕ_0 is the baseline prediction, and $\phi_j^{(t)}$ represents the contribution of feature j (i.e., component j) to the predicted NMH. Since inputs are binary component indicators, each SHAP value shows how including a specific component changes production workload.

Blockage Risk Level Classification

Due to the variability and uncertainty in blockage data, blockage is modeled as a classification problem. Each production scenario is assigned a risk level:

$$r \in \{\text{LOW, MID, HIGH}\} \tag{18.5}$$

A scenario is represented by the following feature vector:

$$\mathbf{z}^{(i)} = \left[\tilde{y}^{(i)}, \text{CPR}^{(i)}, \text{SKUCount}^{(i)}, \text{UniqueParts}^{(i)}, \text{AvgPartsPerSKU}^{(i)}\right]^{\top} \tag{18.6}$$

where $\tilde{y}^{(i)}$ denotes the normalized man-hour value, which is the sum of the NMH values for each SKU and their quantities, used by the classifier (observed NMH when available, and predicted NMH otherwise) and the remaining terms are BOM-derived structural metrics.The classification model is defined as:

$$\hat{r}^{(i)} = G(\mathbf{z}^{(i)}) \tag{18.7}$$

where $G(\cdot)$ assigns each scenario to a risk category.

Common Part Ratio (CPR)

For a product s, let $P(s)$ be the set of components used in that product. The similarity between two products is calculated using Jaccard similarity:

$$J(a, b) = \frac{|P(a) \cap P(b)|}{|P(a) \cup P(b)|} \tag{18.8}$$

For a production scenario with m products, CPR is defined as the average pairwise similarity:

$$CPR(S) = \frac{2}{m(m-1)} \sum_{1 \leq p < q \leq m} J(s_p, s_q) \tag{18.9}$$

By definition, $CPR(S) \in [0, 1]$. Higher CPR indicates more shared components, while lower CPR indicates greater product diversity.

18.4 Validation

Validation was performed across three distinct dimensions: Conceptual Validation (assessing assumptions), Operational Validity on Historical Data (assessing retrospective performance), and Sensitivity Analysis (assessing model continuity). Conceptual validation ensured that the modeling assumptions and structural relationships are aligned with real-world manufacturing logic and expert knowledge. Operational validation evaluated the model's ability to produce consistent and meaningful results when applied to historical production data, confirming its practical relevance. Sensitivity analysis further tested the model's behavior under small input variations, demonstrating stability and continuity in its outputs.

The primary model of OptiCost has an overall accuracy of 0.60. Furthermore, the area under the ROC curve for the LOW-versus-rest boundary is approximately 0.87, which reflects the model's high accuracy especially for the low risk regime. When it comes to the secondary model of OptiCost which predicts the man-hour data for new SKUs, its R-square value of 0.98 reflects the model's consistency. 0.72 hour RMSE of the secondary model indicates that it has less than 1 hour error in its predictions.

Overall, the validation results indicate that the framework reliably captures the relationship between product structure and operational performance, while maintaining robustness and interpretability across different scenarios.

18.5 Benefits to the Company

The complexity department currently lacks sufficient visibility into how production complexity evolves. The proposed tool enables the company to quantify and monitor complexity based on system elements, supporting more informed decision-making.

18.5.1 Blockage Data Classification

To demonstrate the classification model's benefits to the production environment, the model's predictions were utilized in the daily SKU production assignment. The low-risk predictions were prioritized due to their high precision, and the predictions were used as a constraint on the assignment process. This involvement of the model's predictions in the assignment is to reflect the model's capability to reduce the blockage levels with high precision.

For benchmarking purposes, three production days with high blockage, day one with 4762, day two with 2449 and day three with 1135 blocked units, were picked for comparison from the 2025 production data. The SKUs produced and their corresponding quantities were redistributed in a one-month time interval with the restriction of having a low-risk level prediction. The SKUs were redistributed to the days that are considered to be most suitable for them, according to the model inputs with a greedy approach. The SKUs and their quantities that have not satisfied the prediction condition were left unassigned and taken out of production. As a result:

- First day reduced to low risk with 383 units removed.
- Second day reduced to low risk with 513 units removed.
- Third day reduced to low risk with no units removed.

These results demonstrate the model's effectiveness in reducing blockage risk through production reallocation.

18.5.2 BOM-Based Man-Hour Prediction

In this section, normalized man-hour prediction (NMH) model's benefits for the company are going to be demonstrated. The aim is to show that NMH is not only effective in estimating a SKU's average time for production, but also very effective in demonstrating the results of configuration changes applied.

The validation tests on the NMH model showcased high accuracy for estimating a SKU's labor requirements from its BOM structure. With this high accuracy, the impact of new product configurations and engineering (design) changes can be estimated reliably before entering mass production, by the addition of the NMH predictions. The NMH model presents its benefits by reducing the difference between the estimated and actual labor costs, improving the conditions for financial planning and margin protection. For the demonstration of these benefits, benchmark tests were conducted, by choosing an existing product configuration as a baseline and comparing the baseline with its modified versions. Benchmark tests confirm the model's responsiveness to structural changes:

- **Packaging case:** workload increased from 43.6063 to 43.9986 (+0.3922, +0.90%).
- **Water system case:** workload decreased from 47.6929 to 47.5099 (-0.1831, -0.38%).
- **Internal assembly case:** workload increased from 72.6111 to 75.2140 (+2.6029, +3.58%).

Extreme scenarios further validate non-linear behavior:

- **Packaging upgrade: 69.5487 → 65.6709** (-3.8779, -5.58%)
- **Housing upgrade: 43.8701 → 46.0641** (+2.1940, +5.00%)

A continuity test showed high sensitivity to small changes: predicted decrease -0.0177 hours, compared to actual -0.0149 hours (difference: 0.0028). Overall, the model provides a reliable and interpretable tool for evaluating design changes and their operational impact.

18.6 Deliverables

The decision support tool OptiCost has a simple, user-friendly user interface. The screen area used efficiently while designing the user interface, as it is considered as a scarce resource. On top of the screen, there is a navigation bar that allows user to navigate between tabs "Main Panel", "Blockage Analysis", "Man-hour Analysis", "Scenario Analysis", and "Reports". Each tab has different functionality.

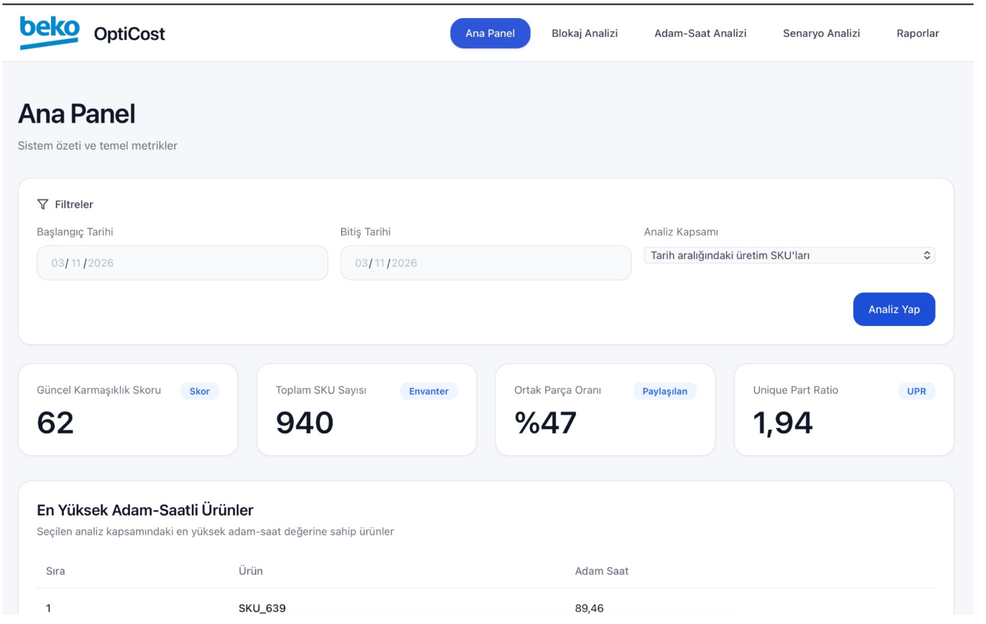

Figure 18.1: Main panel

- **Main Panel:** This tab provides user an overview of the system see Appendix, Figure 18.1. It allows user to filter the analysis period, in other words, user can filter the data to be used in the analysis. Based on the filtered or non-filtered data; complexity score, total SKU count, common part ratio, and unique part ratio are shown by small cards that are located just below the filter options area. Below these cards, there is an area for products that have the highest man-hour requirements. At the end of the screen, 10 most risky parts are listed.

- **Blockage Analysis:** According to the selected date, the blockage risk is predicted in three categories (low, mid, high) by using this tab see Appendix, Figure 18.2.. Also, the previous analyses are located at the end of the screen so user can easily monitor them without conducting the same analysis again.

- **Man-hour Analysis:** By using this tab see Appendix, Figure 18.3, the man-hour requirements for the new product configurations can be predicted accurately. This might be very useful when a new SKU added to the system.

- **Scenario Analysis:** This tab see Appendix, Figure 18.4, allows user to run simulations for selected production plans and selected BOM files. As a result of the simulation, blockage risk for each day throughout the time period that are covered by the production plan displayed.

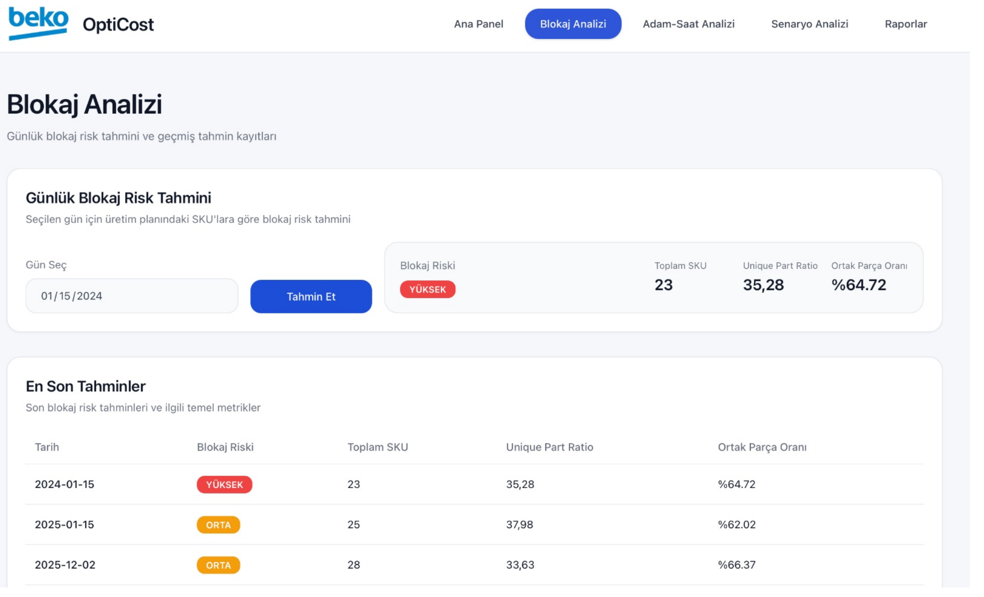

Figure 18.2: Blockage analysis panel

In case of a new SKU added to the system, this tab also shows the predicted man-hour data to the user. At the end of the screen, there is an area that allows user to save the scenario. Then, this saved scenario can be displayed in the "Reports" tab.

- **Reports:** The saved scenarios that are created by using the "Scenario Analysis" tab are listed in this tab. Also, user can download these reports as Excel files for further analysis in Excel. Furthermore, the user can download a comparison report which allows comparing different scenarios easily. OptiCost works on the browser, so it does not require any prior setup process.

18.7 Integration and Implementation

The transition of the OptiCost system from a developmental model to a daily operational tool is structured through a phased implementation plan to ensure seamless adoption by Beko's complexity management team.

- Data Integration and Customization: The Python-based decision-support dashboard is integrated with the factory's existing Bill of Materials (BOM) and historical production datasets. Structural metric calculations (e.g., Common Part Ratio) are carefully calibrated to align with the company's internal reporting standards.

- User Testing and Documentation: The complexity management team

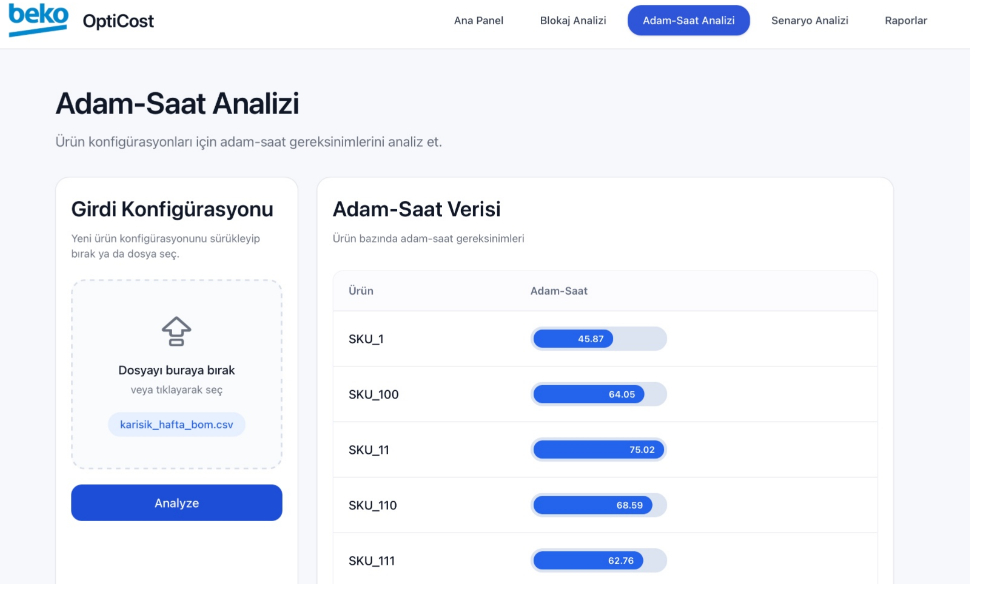

Figure 18.3: Blockage analysis panel

conducts controlled usability testing to evaluate the interface and scenario analysis features. Concurrently, a comprehensive User Manual is developed to ensure the team can independently operate the tool and accurately interpret the SHAP-based explanations.

- Controlled Pilot Study: To mitigate operational risk, initial deployment is conducted as a pilot study on a restricted, representative subset of the product portfolio. Planners evaluate the baseline complexity of this subset against alternative "What-If" scenarios, such as SKU elimination or increased part commonality, to validate the model's real-world accuracy.

- System Refinement and Handover: Feedback gathered during the pilot execution directly informs iterative refinements to the scenario logic, feature calculations, and dashboard UI. The fully validated system and updated documentation are then officially handed over to the company for continuous use.

18.8 Conclusion

The continuous proliferation of product variants and components poses a significant challenge to modern manufacturing efficiency. This project successfully fulfills the gap between structural product design and shop-floor operational realities at the Beko Dishwasher Factory. By using machine

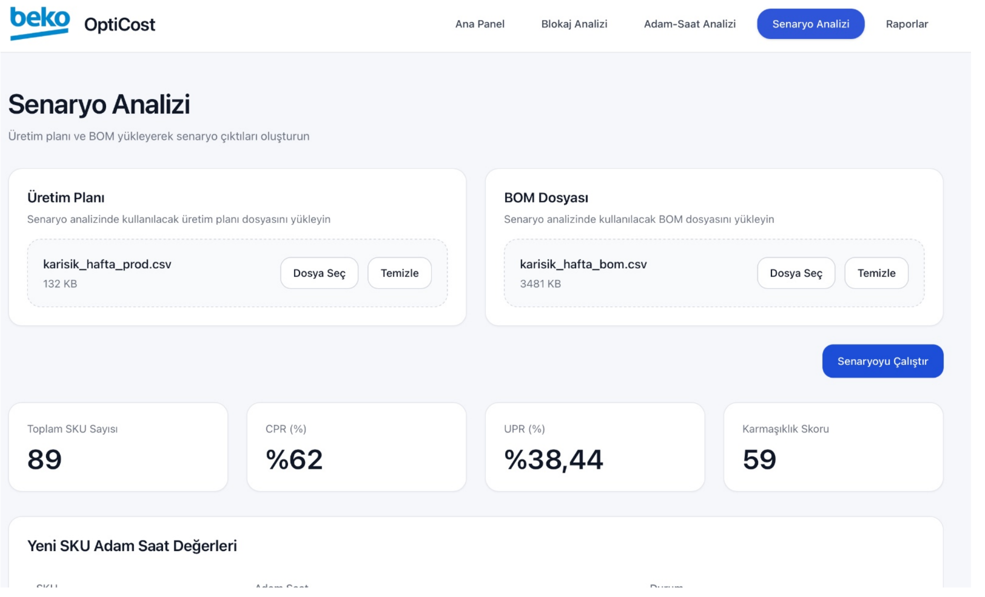

Figure 18.4: Scenerio analysis panel

learning methods, OptiCost transforms system complexity from an unpredictable operational burden into a quantifiable, manageable metric. With its user-friendly and easy-to-use user interface, the decision support tool OptiCost ensures that the company's complexity management team possesses a sustainable, repeatable tool for continuous analysis.

Bibliography

Arçelik Global (2024). 2024 TSRS Uyumlu Sürdürülebilirlik Raporu. https://www.arcelikglobal.com/media/3npp1s5k/2024-tsrs-uyumlu-surdurulebilirlik-raporu.pdf [Online; accessed 26-March-2026].

Beko Corporate (2026). Company Overview — About Us. https://www.bekcorporate.com/company/about-us/overview/ [Online; accessed 26-March-2026].

Bloomberg HT (2024). Arçelik'in Kurumsal Markası Beko Oluyor. https://www.bloomberght.com/arcelik-in-kurumsal-markasi-beko-oluyor-2350608 [Online; accessed 26-March-2026].

Google (2026). Google Earth Pro Software. https://www.google.com/earth/ [Online; accessed 26-March-2026].

Koç Holding (2026). Faaliyet Alanları ve Markalar. https://www.koc.com.tr/faaliyet-alanlari/markalar [Online; accessed 26-March-2026].

Tesis İçi Malzeme Sevkiyatı için Araç Atama ve Rotalama Karar Destek Sistemi | 19

Roketsan

Proje Ekibi
İrem Doğa Demir, Alperen Ezber, Deniz Kutlu,
Öykü Su Merinç, Tuna Öztürk, M. Baha Şahingeri, Aras Tokatlı

Şirket Danışmanı
Bilgehan Bıyıklı
Depo Birim Yöneticisi

Akademik Danışman
Prof. Dr. Oya Karaşan
Endüstri Mühendisliği Bölümü

Özet

Bu projede, Roketsan'ın tesis içi lojistik operasyonlarında teslimat gecikmelerini azaltmak amacıyla araç atama, rotalama ve görev çizelgeleme kararlarını destekleyen bir karar destek sistemi geliştirilmiştir. Problem çok depolu araç rotalama yapısı içinde ele alınmış, büyük ölçekli örneklerde uygulanabilir çözümler elde etmek için Uyarlamalı Büyük Komşuluk Arama sezgisel yöntemi kullanılmıştır. Geliştirilen sistem, günlük rotalama planları üretmenin yanı sıra filo büyüklüğü ile toplam gecikme arasındaki ilişkiyi değerlendirerek teslimat gecikmelerini azaltan daha düzenli bir taşıma planlamasını desteklemektedir.

Anahtar Sözcükler: Tesis içi lojistik, karar destek sistemi, araç rotalama, çok depolu araç rotalama (MDVRP), Uyarlamalı Büyük Komşuluk Arama, Pareto sınırı, gecikme enazlama

A Decision Support System for Vehicle Assignment and Routing in In-Plant Material Transportation

Abstract

In this project, a decision support system was developed to support vehicle assignment, routing, and task scheduling decisions in Roketsan's intralogistics operations with the aim of reducing delivery delays. The problem was addressed within a multi-depot vehicle routing structure, and an Adaptive Large Neighborhood Search heuristic was used to obtain feasible solutions for large-scale instances. In addition to generating daily routing plans, the developed system evaluates the relationship between fleet size and total lateness, thereby supporting more structured transportation planning with reduced delivery delays.

Keywords: Intralogistics, decision support system, vehicle routing, MD-VRP, ALNS, Pareto frontier, delay minimization

19.1 Company Description

Roketsan was established on 14 June 1988, by a decision of the Defense Industry Executive Committee with the aim of meeting the rocket and missile needs of the Turkish Armed Forces (TSK) and establishing a leading institution in the design, development, and production of rockets and missiles in the country. With over 5,000 employees as of 2025, Roketsan is one of the largest companies operating in the defense industry. The company's product portfolio is grouped under land, air defense, naval, precision-guided, space, ballistic production systems, and subsystems. Roketsan operates under strict security and quality standards and has production processes compliant with NATO standards, supported by certifications such as ISO 9001, AS9100, and AQAP-2110 (Roketsan, 2025).

19.2 System Analysis and Problem

19.2.1 System Analysis

Roketsan's intralogistics system operates across the Elmadağ and Lalahan campuses, where raw materials, components, and semi-finished products are transported between warehouses, production units, and testing facilities. These operations include both inter-building and intra-campus movements and play a critical role in maintaining production continuity. The system handles a high volume of transportation requests with varying origins, destinations, and urgency levels. Currently, these operations are carried out

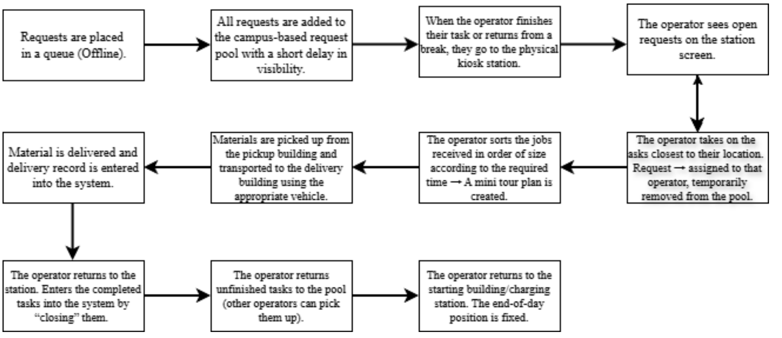

Figure 19.1: Current Process Flow of Material Transport

by a fleet that consists of forklifts, trucks, and commercial vehicles.

The current system does not include real-time vehicle tracking due to strict security regulations, and transportation demand arrives dynamically over time. As a result, the current system follows an operator-based pull structure.

In this system, the transportation requests are generated through an internal system and transferred to a central request pool. As illustrated in Figure 19.1, operators access these requests via kiosks and independently select tasks to execute. Routing and task selection decisions are therefore made locally, without centralized planning. After completing a task, operators return to the kiosks and select new requests.

19.2.2 Problem Definition

The primary objective of the intralogistics system is to ensure timely delivery of materials while maintaining efficient utilization of available resources. However, the current operational structure leads to several inefficiencies that directly affect system performance. Routing and task selection decisions are made locally by operators, which leads to locally efficient but globally suboptimal routes, increasing total travel time and overall system lateness.

Furthermore, transportation requests arrive intermittently and exhibit variability over time. As shown in Figure 19.2, a significant portion of transportation requests have very short time windows, with the majority requiring completion within 0–2 hours. This indicates that the system operates under high time pressure. Under these conditions, the current operational structure leads to inefficient routing and resource utilization, increasing total delivery delays. Therefore, the problem is defined as generating coordinated vehicle assignment, routing, and scheduling decisions that minimize total lateness while satisfying Roketsan's operational constraints. To evaluate system performance, total delivery delay is considered as the primary performance measure, along with the analysis of different fleet size

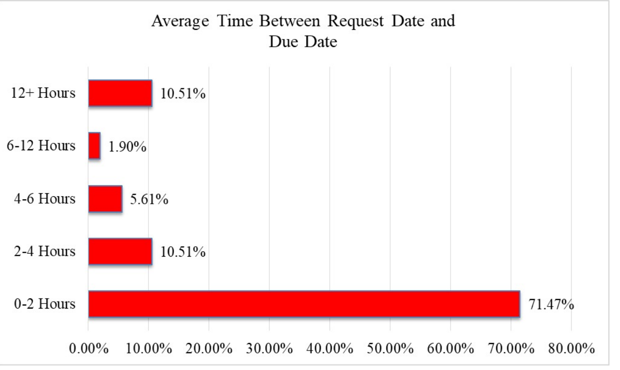

Figure 19.2: Average Time Between Request and Work Completion

scenarios. The system outputs consist of routing and scheduling plans and a set of trade-off solutions presented as a Pareto frontier, allowing planners to compare alternative fleet configurations and their impact on delays.

19.3 Proposed Solution Strategy

19.3.1 Critical Assumptions

The proposed model is based on a set of assumptions to represent the intralogistics operations within Roketsan. Distances between all buildings, as well as between the Elmadağ and Lalahan campuses, are computed on the internal road network using shortest paths. Corresponding travel times for each vehicle type are assumed to be deterministic, and congestion effects are neglected. Trucks and other commercial vehicles are allowed to use all internal and inter-campus roads without restriction. Forklifts, on the other hand, operate under specific constraints. While empty forklifts may move freely, loaded forklifts are restricted from transporting loads over distances exceeding 100 meters. In addition, forklifts are not allowed to travel directly between the Elmadağ and Lalahan campuses. Any long-distance repositioning of forklifts is therefore assumed to be performed by loading them onto trucks or other suitable carriers. Finally, the model developed within the scope of this project is designed to be executed on a daily basis, generating routing and scheduling plans within the required planning horizon.

19.3.2 Major Constraints

The constraints considered in the model are categorized into two main groups: model-based constraints and company regulations. Model-based constraints ensure the feasibility and consistency of routing and scheduling decisions. These constraints include that each vehicle starts and ends its route at a designated depot within the Elmadağ or Lalahan campuses. Flow conservation constraints maintain route continuity, while each pickup–delivery job is served exactly once by a single vehicle with precedence between pickup and delivery. Vehicle capacity constraints are enforced, and handling times are explicitly incorporated into routing decisions. In addition, the number of forklifts available for handling is limited by their prior allocation to transportation tasks, ensuring that total resource usage does not exceed the available fleet. Company regulations impose additional operational limitations. Due to strict information security policies, technologies such as GPS cannot be used, requiring planning without real-time connectivity (Bıyıklı, 2025). Vehicle movements must also comply with internal safety and regulations, including restrictions on accessible areas and load types.

19.3.3 Objective

The proposed model's objective is tominimize the total lateness of pickup–delivery jobs in Roketsan's intralogistics system, where lateness is defined as the positive deviation of a job's completion time from its due date. By focusing on lateness rather than travel distance, the model directly targets the timely completion of transportation tasks critical for production continuity. In addition, the model is solved under different fleet-size limits within an ε-constraint framework to capture the trade-off between delivery performance and vehicle utilization (Mavrotas, 2009). The resulting transportation plan is complemented by a second-stage model that aims to minimize the number of forklifts deployed from the depot to satisfy local handling requirements.

19.3.4 The Conceptual Model

The proposed decision support system consists of two sequential optimization modules based on deterministic demand data. For each planning horizon, transportation requests, vehicle information, and network data are converted into standardized job and distance matrices.

The transportation problem is formulated as a multi-depot vehicle routing problem (MDVRP) and solved using an ε-constraint approach with an ALNS-based heuristic. In this module, the system generates vehicle assignment, routing, and scheduling decisions for transportation requests across campuses. Alternative solutions are obtained under different fleet-size limits to evaluate the trade-off between delivery performance and resource usage. In the handling module, the resulting transportation plan is used to model local handling operations at each campus. Based on the timing outputs, the required forklift capacity is determined to support loading and unloading activities.

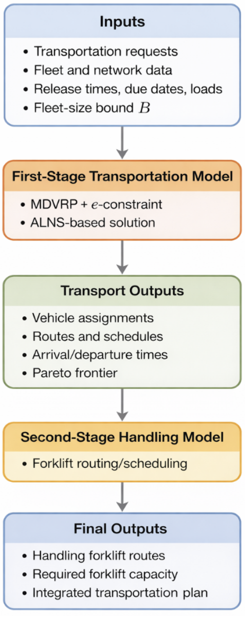

Figure 19.3: Flow Chart of the Conceptual Model

19.3.5 Mathematical Model

The transportation model is formulated as a mixed-integer linear programming (MILP) model for a multi-depot vehicle routing problem defined on a network representing both campuses. The mathematical model includes sets, parameters, and decision variables that capture transportation requests, vehicle types, travel times, and constraints. Each request is represented as a pickup–delivery job with associated release times, due dates, loads, and handling durations. In the subsequent handling model, the timing outputs obtained from the transportation module are used to model local handling operations. A separate MILP formulation is constructed for each campus, where handling tasks are represented as time-dependent jobs. This model determines the number of forklifts required to execute loading and unloading operations within the given time windows. The detailed formulations for each model, including sets, parameters, decision variables, and constraints of both stages, are provided in Appendix 19.A and Appendix 19.B.

19.3.6 Heuristic Solution Approach

The first-stage problem is formulated as a multi-depot pickup–delivery MILP, which is NP-hard. Since realistic daily instances are too large to be solved efficiently by exact methods, an Adaptive Large Neighborhood Search (ALNS) heuristic is used to obtain high-quality solutions within practical computation times (Ropke and Pisinger, 2006). ALNS is well suited to large-scale and tightly constrained vehicle routing problems because it iteratively improves a solution through destroy–repair steps while adaptively favoring the operators that perform better over time. For a given fleet-size bound, the algorithm first generates an initial feasible solution and then improves it by removing and reinserting subsets of jobs in promising positions. The objective of the heuristic is to minimize total lateness:

$$f(s) = \sum_{p \in P} \max\{0,\, C_p(s) - D_p\},$$

where s denotes a feasible routing solution, P is the set of transportation jobs, $C_p(s)$ is the completion time of job p under solution s, and D_p is the due date of job p. In other words, the heuristic evaluates each solution according to the total delay of all jobs completed after their due dates. A simulated annealing-based acceptance rule is used to allow controlled diversification during the search process. The heuristic was tested against the exact model on smaller instances and produced solutions with small optimality gaps while preserving the expected trade-off between fleet size and total lateness.

19.4 Validation

The validation stage aimed to assess whether the proposed framework provides a credible representation of Roketsan's real intralogistics operations. Although the overall solution structure was designed for a two-campus system, validation was conducted only for the Elmadağ campus due to data availability limitations. In particular, the detailed inter-building distance matrix for the Lalahan campus was not available, so the validation scope

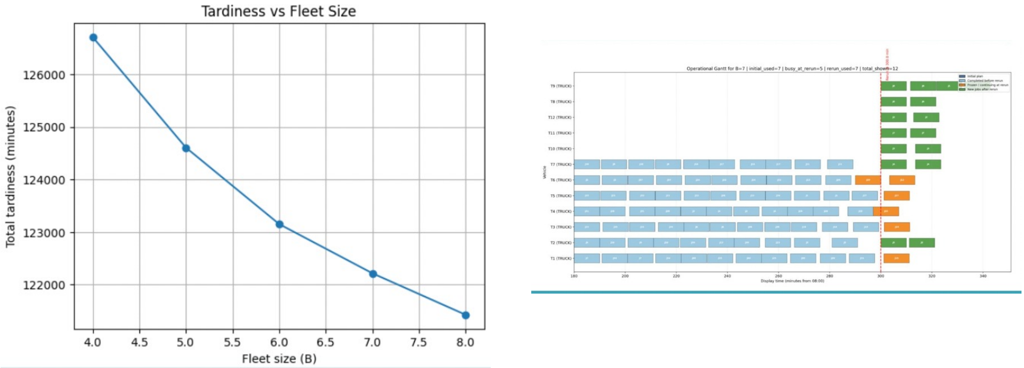

Figure 19.4: Pareto frontier and vehicle Gantt chart for the validation instance dated 28.11.2025.

was restricted to Elmadağ. The validation experiments were carried out using real company data, including the Elmadağ distance matrix and historical transportation requests recorded between October 2024 and December 2025. These records contained origin and destination building information, request dates, due dates, and package quantities. In addition, the modeling assumptions and parameter settings used in the study were reviewed and confirmed with company representatives.

Three representative scenarios were analyzed under different fleet-size bounds. The first scenario, based on 01.12.2025, included 97 jobs and showed a clear reduction in total tardiness as fleet size increased. A re-run interval of 2 hours was used, to oversee demand in 2-hour time intervals. For this date, 18 forklifts satisfied all demand for an upper fleet size of 8. The second scenario, representing a peak-load day on 28.11.2025 with 142 jobs, exhibited the same pattern under heavier demand, confirming the model maintains consistent performance under congested conditions; see Figure 19.4.

The third scenario, based on 29.11.2025 with 36 jobs, represented a low-demand day and showed that beyond a certain fleet size, additional vehicles did not yield meaningful improvement. In addition to tardiness results, the generated routes and schedules were evaluated. The solutions followed coherent building sequences, avoided unnecessary back-and-forth movements, and satisfied vehicle-type restrictions, pickup-delivery precedence, and depot constraints. Overall, the validation findings supported the use of the proposed system as a realistic and practically applicable decision support tool.

19.5 Decision Support System Interface

A user-oriented interface was developed to support the practical use of the decision support system. Through the interface, users can upload daily operational data and define planning parameters such as fleet size and rerun intervals. This allows users to review the transportation plan and use the generated outputs in daily planning decisions. Example decision support system screens can be seen in Figure 19.5. After the system is executed, the outputs presented in the DSS to the company are the used vehicles, their routes in the distance matrix, the Pareto chart of tardiness vs. fleet upper bound, and the Gantt charts of both transportation vehicles and forklifts.

19.6 Implementation and On-Site Evaluation

The proposed decision support system was developed and demonstrated with company input to assess its practical applicability in Roketsan's intral-

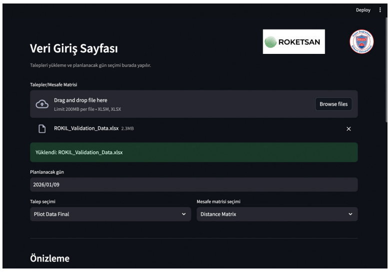

Decision Support System Execution Page

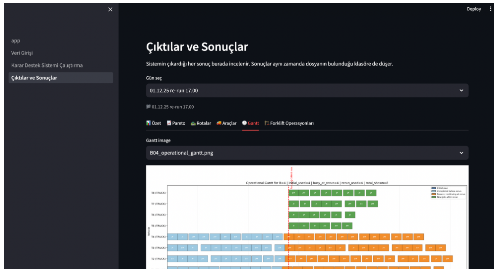

Results and Outputs Page

Figure 19.5: User Interface Screens of the Decision Support System

ogistics operations. Based on the evaluations conducted during the project, the model, data preparation steps, and user interface were refined to improve consistency and usability. The resulting system provides a structured tool for daily routing and scheduling decisions. Its practical value is expected to increase further as data integration improves and the company's logistics management system becomes more fully incorporated into the planning process.

19.7 Benefits to the Company

The proposed decision support system provides measurable improvements in Roketsan's intralogistics performance, particularly in terms of total tardiness, which was selected as the main key performance indicator. Benchmarking against the current system was conducted using historical daily instances from December 2025, and the proposed framework was evaluated under different fleet-size bounds. The resulting performance improvements, expressed as percentage reductions in total tardiness compared to current operations, are summarized in Table 19.1. The results indicate that the proposed framework generates consistent operational gains under different workload levels. The timespan of 12-22 December was used, since the days indicated different demand levels on different weekdays. The highest improvement was observed on 13.12.2025, reaching up to 16.38%. This day seperates from the others. The moderate results, closer to the average, are observed on 16.12.2025, with the improvement of 6.893% and 18.12.2025, with 6.896%. Across the selected timeline, the average improvement was calculated as 6.769%, indicating a consistent reduction in tardiness under different demand levels. In addition to reducing total tardiness, the system enables planners to observe how performance varies with fleet size and to identify capacity levels beyond which additional vehicles provide limited marginal benefit. Therefore, the project contributes not only as a routing and scheduling tool, but also as a managerial decision support system for evaluating resource utilization and service performance.

Date	Weekday	Number of Jobs	B	Lateness	DSS Lateness	Improvement (%)
12.12.2025	5	93	4	80865	75931	6.102
13.12.2025	6	9	1	7100	5937	16.38
15.12.2025	1	77	4	68372	62309	8.868
16.12.2025	2	115	4	109747.32	102182.65	6.893
17.12.2025	3	133	4	124272.26	122724.00	1.246
18.12.2025	4	45	3	38216.35	35581.00	6.896
19.12.2025	5	79	4	68374.36	62806.66	8.143
20.12.2025	6	18	1	10865.04	10564.00	2.771
21.12.2025	7	8	1	5914.5	5553.8	6.099
22.12.2025	1	82	4	68568.18	65625.18	4.292
Average	-	-	-	-	-	**6.769**

Table 19.1: Weekly tardiness and improvement rates

19.8 Conclusion and Recommendations

This project successfully addresses Roketsan's need for a structured and centralized decision support system in intralogistics operations, where routing and task assignment were previously performed in a decentralized and operator-based manner. The proposed two-stage framework provides coordinated vehicle assignment, routing, and scheduling decisions, directly targeting the reduction of total tardiness, which was identified as the primary performance metric. Validation results based on real operational data demonstrate that the system consistently achieves significant improvements, with consistent reductions in total tardiness across different daily instances, indicating that the project meets the organization's expectations in improving service performance and operational efficiency. In addition, the ability to evaluate fleet-size trade-offs through Pareto analysis supports managerial decision-making beyond daily planning. For future work, it is recommended to extend the on-site evaluation into a broader implementation phase based on operational feedback. Extending the model to fully incorporate both campuses, improving data integration processes, and exploring the use of dynamic or real-time information—subject to security constraints—would further enhance the applicability of the system. Overall, the proposed approach provides a scalable and practical solution that can support both operational and strategic decision-making in intralogistics systems.

Bibliography

Bıyıklı, B. (2025). Personal interview. Elmadağ, Turkey: Roketsan. Oct 8.

Mavrotas, G. (2009). Effective implementation of the ϵ-constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation 213*(2), 455–465.

Roketsan (2025). About us. Retrieved October 15.

Ropke, S. and D. Pisinger (2006). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science 40*(4), 455–472.

Appendix: Mathematical Model

19.A First-Stage Model

Table 19.2: First-Stage Model Sets and Parameters

Notation	**Description**
$D = \{n_E, n_L\}$	Set of depots (Elmadağ and Lalahan).
N	Set of all nodes, where $D \subseteq N$.
P	Set of pickup–delivery jobs.
K_T	Set of trucks.
K_F	Set of forklifts.
$K = K_T \cup K_F$	Set of all vehicles.
$A^{>100}$	Set of arcs with distance greater than 100 meters.
$C = N \setminus D$	Set of customer nodes.
o_p	Pickup node of job $p \in P$.
d_p	Delivery node of job $p \in P$.
t_{ij}	Travel time from node i to node j.
d_{ij}	Distance between nodes i and j.
q_p	Load of job $p \in P$.
Q_k	Capacity of vehicle $k \in K$.
r_p	Release time of job $p \in P$.
D_p	Due date of job $p \in P$.
$a_{p,\text{pick}}$	Handling time for pickup of job p.
$a_{p,\text{del}}$	Handling time for delivery of job p.
M	Big-M constant.
B	Fleet size bound used in the ε-constraint method.

Decision Variables

- $x_{ijk} \in \{0, 1\}$: Equals 1 if vehicle k travels from node i to node j.
- $y^p_{ijk} \in \{0, 1\}$: Equals 1 if job p is carried by vehicle k on arc (i, j).
- $T_{ik} \geq 0$: Arrival time of vehicle k at node i.

- $G_{ik} \geq 0$: Departure time of vehicle k from node i.
- $z_k \in \{0,1\}$: Vehicle usage indicator.
- $L_p \geq 0$: Lateness of job p.

Objective Function

$$\min \quad \sum_{p \in P} L_p, \tag{19.1}$$

s.t.

$$\sum_{k \in K} z_k \leq B, \tag{19.2}$$

$$\sum_{d \in D} \sum_{j \in N} x_{djk} = z_k \quad \forall k \in K, \tag{19.3}$$

$$\sum_{d \in D} \sum_{i \in N} x_{idk} = z_k \quad \forall k \in K, \tag{19.4}$$

$$\sum_{j \in N} x_{djk} = \sum_{i \in N} x_{idk} \quad \forall d \in D \quad \forall k \in K, \tag{19.5}$$

$$x_{ijk} \leq z_k \quad \forall i, j \in N \quad \forall k \in K, \tag{19.6}$$

$$\sum_{k \in K} \sum_{j \in N} y^p_{o_p jk} = 1 \quad \forall p \in P, \tag{19.7}$$

$$\sum_{k \in K} \sum_{j \in N} y^p_{j d_p k} = 1 \quad \forall p \in P, \tag{19.8}$$

$$\sum_{j \in N} x_{ijk} = \sum_{j \in N} x_{jik} \quad \forall i \in N \quad \forall k \in K, \tag{19.9}$$

$$y^p_{ijk} \leq x_{ijk} \quad \forall p \in P \quad \forall i, j \in N \quad \forall k \in K, \tag{19.10}$$

$$\sum_{j \in N} y^p_{ijk} = \sum_{j \in N} y^p_{jik} \quad \forall p \in P \quad \forall i \in C \quad \forall k \in K, \tag{19.11}$$

$$\sum_{p \in P} q_p y^p_{ijk} \leq Q_k \quad \forall i, j \in N \quad \forall k \in K, \tag{19.12}$$

$$G_{ik} \geq T_{ik} + \sum_{\substack{p \in P \\ o_p = i}} a_{p,\text{pick}} \sum_{j \in N} y^p_{ijk} + \sum_{\substack{p \in P \\ d_p = i}} a_{p,\text{del}} \sum_{j \in N} y^p_{ijk} \quad \forall i \in N \quad \forall k \in K, \tag{19.13}$$

$$T_{jk} \geq G_{ik} + t_{ij} - M(1 - x_{ijk}) \quad \forall i \in N \quad \forall j \in C \quad \forall k \in K, \tag{19.14}$$

$$T_{o_p k} \geq r_p - M\Big(1 - \sum_{j \in N} y^p_{o_p jk}\Big) \quad \forall p \in P \quad \forall k \in K, \tag{19.15}$$

$$G_{d_p k} \leq D_p + L_p + M\Big(1 - \sum_{j \in N} y^p_{j d_p k}\Big) \quad \forall p \in P \quad \forall k \in K, \tag{19.16}$$

$$T_{d_p k} \geq G_{o_p k} - M\left(1 - \sum_{j \in N} y^p_{o_p jk}\right) \quad \forall p \in P \quad \forall k \in K, \tag{19.17}$$

$$\sum_{\substack{j \in N \\ j \neq o_p}} y^p_{o_p jk} = \sum_{\substack{i \in N \\ i \neq d_p}} y^p_{id_p k} \quad \forall p \in P \quad \forall k \in K, \tag{19.18}$$

$$\sum_{p \in P} y^p_{ijk} = 0 \quad \forall (i,j) \in A^{>100} \quad \forall k \in K_F, \tag{19.19}$$

$$x_{ijk},\, y^p_{ijk},\, z_k \in \{0,1\} \quad \forall p \in P \quad \forall i,j \in N \quad \forall k \in K, \tag{19.20}$$

$$T_{ik},\, G_{ik},\, L_p \geq 0 \quad \forall i \in N \quad \forall k \in K \quad \forall p \in P. \tag{19.21}$$

19.B Second-Stage Model

Table 19.3: Second-Stage Model Sets and Parameters

Notation	Description
$J = \{1, \ldots, n\}$	Set of handling jobs obtained from the first-stage solution.
0	Depot node.
s_i	Start time of job $i \in J$.
f_i	Finish time of job $i \in J$.
t_{ij}	Travel time from job i to job j, $\forall i, j \in J$.
t_{0i}	Travel time from depot to job i.
A	Set of feasible arcs, where $A = \{(i,j) \in J \times J \mid f_i + t_{ij} \leq s_j\}$.
M	Big-M constant.

Decision Variables

- $x_{ij} \in \{0,1\}$: Equals 1 if a forklift moves from job i to job j.
- $x_{0i} \in \{0,1\}$: Equals 1 if a forklift starts its route with job i.
- $x_{i0} \in \{0,1\}$: Equals 1 if a forklift returns to the depot after job i.
- $a_i \geq 0$: Arrival time at the location of job i.

Objective

$$\begin{aligned}
\min \quad & \sum_{i \in J} x_{0i} \\
\text{s.t.} \quad & \\
& x_{0i} + \sum_{j:(j,i) \in A} x_{ji} = 1 \quad \forall i \in J, \\
& x_{i0} + \sum_{j:(i,j) \in A} x_{ij} = 1 \quad \forall i \in J, \\
& \sum_{i \in J} x_{0i} = \sum_{i \in J} x_{i0},
\end{aligned}$$

$$a_i \leq s_i \quad \forall i \in J,$$
$$a_j \geq f_i + t_{ij} - M(1 - x_{ij}) \quad \forall (i,j) \in A,$$
$$a_i \geq t_{0i} x_{0i} \quad \forall i \in J,$$
$$x_{ij} \in \{0,1\} \quad \forall (i,j) \in A,$$
$$x_{0i},\, x_{i0} \in \{0,1\} \quad \forall i \in J,$$
$$a_i \geq 0 \quad \forall i \in J.$$

Toz Deterjan Paketleme Hatlarında Çizelgeleme İyileştirmesi | 20

Unilever Ev ve Kişisel Bakım Ürünleri Fabrikası

Proje Ekibi

Betül Aslan, Berfin Çeri, Zeynep Erva Dama, Gökhan Gülçakır, Aziz Furkan Söğütlü, Simay Sökmen, Egenaz Tunçdemir

Şirket Danışmanı
Bilgehan Tamkoç
Toz Üretim ve BiH TR Projesi
Türkiye Kıdemli Müdürü

Akademik Danışman
Doç. Dr. Emre Nadar
Endüstri Mühendisliği Bölümü

Özet

Fabrikanın toz deterjan paketleme süreci için karar destek sistemi tasarlanmıştır. Sekiz paralel paketleme hatlı fabrikada, haftalık paketleme planlaması uzman görüş yardımıyla yürütülmektedir. Bu çalışmada, aynı sürecin sistematik bir çizelgeleme yaklaşımıyla iyileştirilmesi hedeflenmiştir. İki aşamalı sözlüksel model, öncelik kuralları, hat bakım programları ve hat öncesi üretim kısıtlarını gözeterek ilk aşamada sistemin tamamlanma süresini, ikinci aşamada hatların toplam tamamlanma süresini azaltmayı hedeflemiştir. Çözüm sürelerini makul düzeyde tutmak için açgözlü sezgisel yöntem başlangıç çözümünü oluşturmakta ve Tabu Arama algoritması çözümü iyileştirmek için kullanılmaktadır. Karar destek sistemi hat tamamlanma süresinde %2,42 oranında, bütün hatların toplam çalışma süresinde %10,26 oranında iyileştirme sağlamaktadır.

Anahtar Sözcükler: Paketleme, karar destek sistemi, paralel hat çizelgelemesi, sözlüksel model, toplam paketleme süresi, Tabu Arama

Improving Packaging Line Scheduling for Powder Detergents

Abstract

This project presents the design and implementation of a decision support system (DSS) for the powder detergent packaging process at the Unilever Konya Home Care Factory. The facility operates eight parallel lines packing various SKUs, and weekly packaging schedules are currently manually generated based on expert judgment. This project aims to improve this process through a systematic scheduling approach. The developed two-stage lexicographic model aims to reduce the overall completion time (makespan) in the first stage and improve the sum of line completion times in the second stage, while incorporating priority rules, maintenance schedules, and upstream production constraints. Considering the solution times required for practical application, a hybrid heuristic approach is adopted, where a greedy heuristic is used to construct the initial schedule and a Tabu Search algorithm is employed for solution improvement. The proposed DSS yields a 2.42% improvement in makespan and a 10.26% reduction in sum of line completion times.

Keywords: Packaging, decision support system, parallel line scheduling, lexicographic model, makespan, Tabu Search

20.1 System Analysis and Problem

The Unilever Konya Home Care Factory is the company's largest powder detergent production site in Europe, producing approximately 600,000 tons per year and supplying 39 countries across Europe, the Middle East, and Africa. The facility operates continuously on a 24/7 basis, with highly automated lines and a small workforce mainly focused on maintenance, changeovers, and quality control.

The powder detergent unit, accounting for over one-third of the factory's capacity, produces about 200,000 tons annually across 20 formulations corresponding to more than 60 SKUs. While production and post-processing operate continuously, the packaging department is both the final stage and the main bottleneck of the facility. Eight parallel packaging lines must frequently switch between SKUs and formats, requiring cleaning and adjustments that can take 15 minutes to 6 hours. Consequently, the sequence of SKUs directly affects line utilization, sum of line completion times and overall factory throughput.

Currently, weekly packaging schedules are prepared manually based on intermediate product levels, base powder availability, and SKU demand.

These schedules rely heavily on expert judgment; consequently, they remain largely static and difficult to adapt in the face of operational disruptions. Limited real-time visibility of inventories further constrains scheduling flexibility, often leading to manual adjustments and extended completion times on the shop floor.

20.2 Proposed Solution Strategy

The company aims to minimize the total duration of weekly packaging operations by generating schedules that reduce line completion times. Historical schedules were analyzed to characterize existing operational practices and to inform the development of the mathematical model presented in Section 20.2.2.

20.2.1 Critical Assumptions and Major Constraints

The problem setting assumes stable upstream supply, fixed weekly demand, and predetermined base plans, with unforeseen disruptions—such as machine breakdowns or labor absences—managed by rerunning the DSS. The scope encompasses eight packaging lines, each processing a single product at a time and requiring cleaning and setup for changeovers. Key operational constraints include limited simultaneous operations due to shared operators and the requirement that priority products be completed early in the week. The mathematical model in the following section is developed based on this operational setting.

20.2.2 Mathematical Model

This subsection presents the lexicographic mixed-integer linear programming (MILP) approach adopted for the weekly scheduling problem. The problem is solved in two sequential stages using MILP formulations: Stage 1 balances the workload across lines by minimizing the makespan (C_{max}), ensuring the shortest possible duration for the entire weekly production. Stage 2 then minimizes the sum of line completion times ($\sum_{\ell \in L} C^{\ell}$) while using the optimal makespan from Stage 1 as a strict upper bound. This hierarchical structure captures both operational efficiency and workload distribution (Pinedo, 2022).

The model accounts for sequence-dependent changeovers and line capacities, and it enforces strict precedence for priority products, while the remaining products are sequenced to optimize overall efficiency.

Notation (Common for Both Stages)

Sets: J: SKU orders; L: Production lines; $L_{sync} \subset L$: Synchronized lines (Lines 2, 3); $J_\ell = J \cup \{s_\ell, f_\ell\}$: Job set for line ℓ; R: Priority jobs.

Parameters:

$$p_{j\ell} : \text{Processing time of job } j \text{ on line } \ell,$$
$$s_{ij\ell} : \text{Sequence-dependent changeover time } (i \to j) \text{ on line } \ell,$$
$$e_{j\ell} = \begin{cases} 1, & \text{if job } j \text{ is eligible for line } \ell, \\ 0, & \text{otherwise,} \end{cases}$$
$$A_{\ell} = \begin{cases} 1, & \text{if line } \ell \text{ is available,} \\ 0, & \text{otherwise,} \end{cases}$$
$$\tau_j = \begin{cases} 3, & \text{machine delay (intermediate product available),} \\ 4, & \text{post-dose delay (base powder only),} \\ 8, & \text{tower preparation delay (no material available).} \end{cases}$$
$$\delta : \text{Maximum completion time difference (Line 5 vs Line 6),}$$
$$M : \text{Sufficiently large constant (Big-}M\text{).}$$

Decision Variables:

$$x_{j\ell} = \begin{cases} 1, & \text{if job } j \text{ is assigned to line } \ell, \\ 0, & \text{otherwise,} \end{cases}$$
$$y_{ij\ell} = \begin{cases} 1, & \text{if job } j \text{ follows job } i \text{ on line } \ell, \\ 0, & \text{otherwise,} \end{cases}$$
$$z_{ij} = \begin{cases} 1, & \text{if job } i \text{ precedes job } j \text{ (both on } L_{sync}\text{),} \\ 0, & \text{otherwise,} \end{cases} \qquad (\forall i, j \in J, i < j)$$
$$t_j \geq 0 : \text{Start time of job } j,$$
$$c_j \geq 0 : \text{Completion time of job } j,$$
$$C^{\ell} \geq 0 : \text{Completion time of line } \ell,$$
$$C_{\max} \geq 0 : \text{Makespan (maximum line completion time).}$$

Stage 1: Minimizing the Makespan

$$\min C_{\max}$$

Constraints:

Assignment Feasibility Constraints

$$x_{j\ell} \leq e_{j\ell} A_{\ell}, \qquad \forall j \in J, \forall \ell \in L \tag{20.1}$$
$$\sum_{\ell \in L} x_{j\ell} = 1, \qquad \forall j \in J \tag{20.2}$$

Flow Conservation and Sequencing with Dummy Nodes Constraints

$$\sum_{i \in J_\ell \setminus \{j\}} y_{ij\ell} = x_{j\ell}, \qquad \forall j \in J, \forall \ell \in L \tag{20.3}$$

$$\sum_{k \in J_\ell \setminus \{j\}} y_{jk\ell} = x_{j\ell}, \qquad \forall j \in J, \forall \ell \in L \tag{20.4}$$

$$\sum_{i \in J_\ell \setminus \{s_\ell\}} y_{is_\ell\ell} = 0, \qquad \forall \ell \in L \tag{20.5}$$

$$\sum_{j \in J_\ell \setminus \{f_\ell\}} y_{f_\ell j\ell} = 0, \qquad \forall \ell \in L \tag{20.6}$$

$$\sum_{j \in J} y_{s_\ell j\ell} \leq 1, \qquad \forall \ell \in L \tag{20.7}$$

$$\sum_{i \in J} y_{if_\ell\ell} \leq 1, \qquad \forall \ell \in L \tag{20.8}$$

Priority Enforcement and Temporal Constraints

$$t_i \geq t_j + p_{j\ell} - M(2 - x_{i\ell} - x_{j\ell}), \qquad \forall j \in R,\ \forall i \in J \setminus R,\ \forall \ell \in L \tag{20.9}$$

$$t_j \geq t_i + p_{i\ell} + s_{ij\ell} - M(1 - y_{ij\ell}), \qquad \forall i, j \in J,\ i \neq j,\ \forall \ell \in L \tag{20.10}$$

$$c_j \geq t_j + p_{j\ell} - M(1 - x_{j\ell}), \qquad \forall j \in J, \forall \ell \in L \tag{20.11}$$

Synchronization and Inventory-Aware Constraints

$$t_j \geq c_i - M(1 - z_{ij}) - M\left(2 - \sum_{\ell \in L_{sync}} x_{i\ell} - \sum_{\ell \in L_{sync}} x_{j\ell}\right), \quad \forall i, j \in J, i < j \tag{20.12}$$

$$t_j \geq \tau_j - M\Big(1 - \sum_{\ell \in L} y_{s_\ell j\ell}\Big), \qquad \forall j \in J \tag{20.13}$$

Line Completion and Workload Balancing

$$C^\ell \geq c_j - M(1 - x_{j\ell}), \qquad \forall j \in J,\ \forall \ell \in L \tag{20.14}$$

$$\delta \geq C^5 - C^6 \tag{20.15}$$

$$\delta \geq C^6 - C^5 \tag{20.16}$$

Stock Control

$$t_j \ \geq \ \tau_j \ - \ M\Big(1 - \sum_{\ell \in L} y_{s_\ell j\ell}\Big) \qquad \forall j \in J \tag{20.17}$$

Makespan Bounding Constraint

$$C_{\max} \geq C^\ell, \qquad \forall \ell \in L \tag{20.18}$$

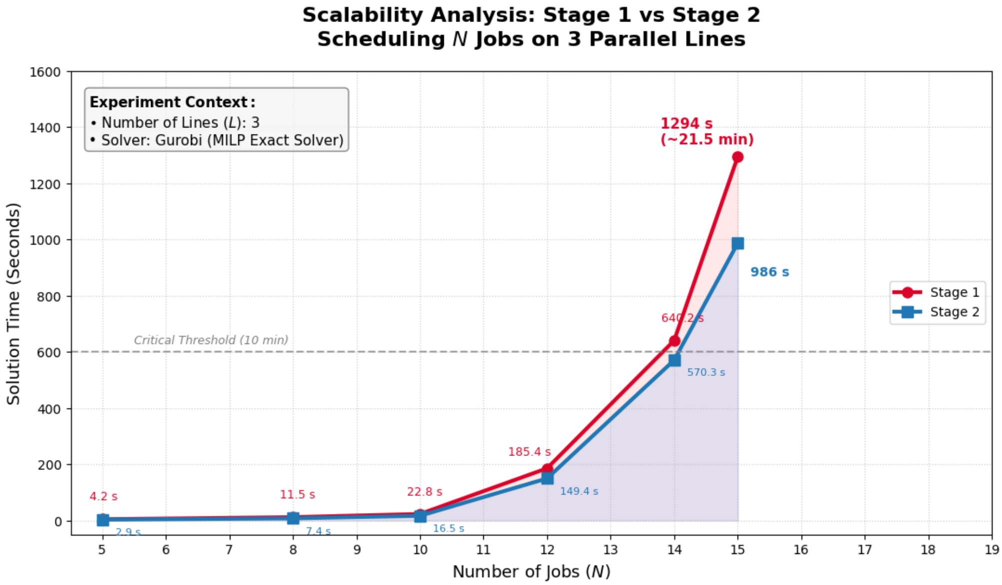

Figure 20.1: Scalability analysis of the MILP formulation on three parallel lines using the Gurobi solver.

Stage 2: Minimizing the Sum of Line Completion Times

$$\min \sum_{\ell \in L} C^{\ell}$$

Additional Constraint: Alongside constraints 20.1–20.18, the following constraint is introduced to utilize Stage 1 optimal solution as an upper bound.

$$C^{\ell} \leq C^{*}_{\max}, \qquad \forall \ell \in L$$

20.2.3 Heuristic Solution Method

As shown in Figure 20.1, solution times increase exponentially with the number of jobs. Consequently, a hybrid heuristic framework is adopted to generate high-quality schedules within practical time limits, ensuring reliable decision support in dynamic operational settings.

The heuristic focuses on generating detailed sequencing and line assignment decisions within a fixed weekly horizon, consistent with operational planning requirements. The approach mirrors the two-stage structure of the MILP by employing a constructive initialization phase followed by an improvement phase utilizing Tabu Search.

Initialization Phase: An initial schedule is constructed using a greedy heuristic: jobs are sorted by priority and assigned to eligible lines that

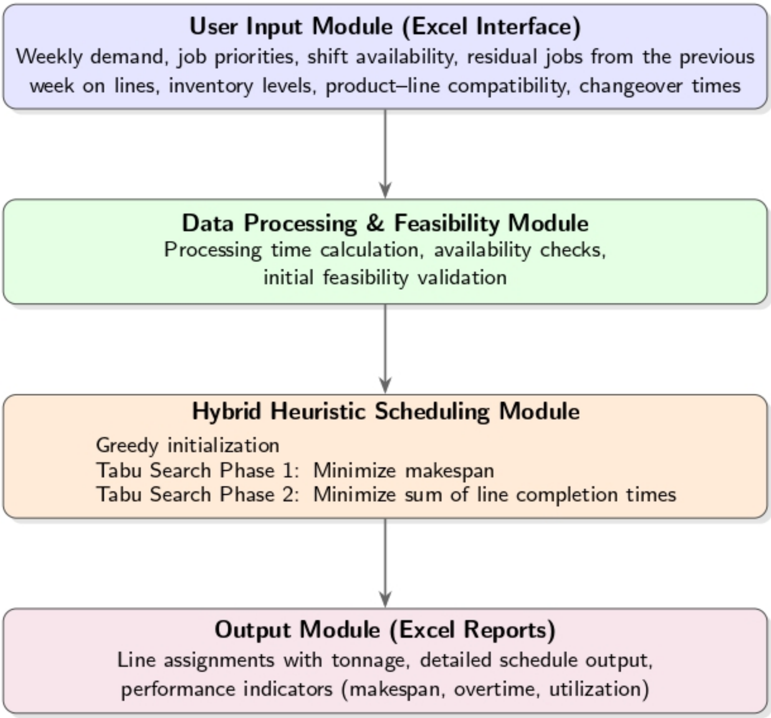

Figure 20.2: Conceptual Model Showing the Major Components, Inputs, and Outputs of the Proposed Scheduling System

yield the earliest completion time. This process accounts for shift availability, sequence-dependent changeovers, and inventory constraints, providing a feasible starting point for the subsequent improvement phase.

Tabu Search Improvement Phase: Tabu Search is a robust metaheuristic widely used for complex scheduling problems due to its ability to explore large solution spaces and escape local optima (Glover, 1989, 1990). Starting from the initial solution, the algorithm explores neighboring schedules through job relocations and swaps while maintaining feasibility. Tabu memory structures are utilized to prevent cycling by restricting recently visited moves.

Consistent with the MILP structure, the heuristic operates in two stages: Stage 1 focuses on reducing the makespan (C_{max}). Stage 2 primarily minimizes the sum of line completion times ($\sum_{\ell \in L} C_\ell$), using the best Stage 1 value as a reference. Unlike the MILP structure, this heuristic framework allows the Stage 1 makespan to be further refined during the second stage if a better solution is identified.

The solution is structured as a Lexicographic Heuristic that transforms

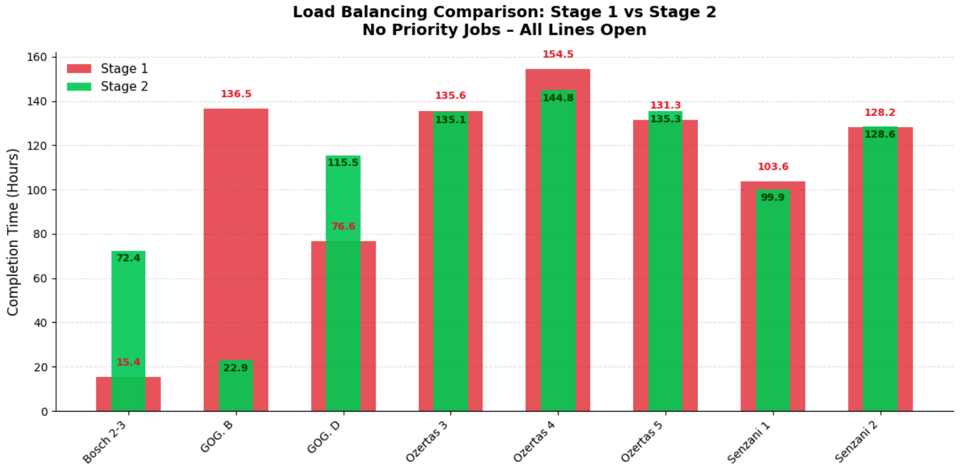

Figure 20.3: Stage 1 and Stage 2 schedule comparison.

operational inputs into a feasible weekly schedule, as illustrated in Figure 20.2. Following initial data processing, the hybrid framework applies the greedy initialization and the two-stage Tabu Search. The final output is delivered via an Excel interface, providing detailed line assignments and key performance indicators.

20.3 Solution Performance and Validation

The proposed scheduling framework was evaluated through both controlled experiments and real production data to ensure correctness and practical effectiveness.

Model Soundness: A verification phase was conducted to confirm the logical integrity of the mathematical model and the heuristics, ensuring their outputs aligned with expected theoretical behaviors. Synthetic instances were used to verify that sequence-dependent changeovers, line eligibility constraints, and priority rules were correctly enforced. The experiments also indicated that the selected lexicographic structure, which prioritizes $C_{\max}$ before $\sum_{\ell \in L} C^{\ell}$, produced a more balanced workload distribution and lower overtime than alternative objective orderings.

Solution Performance: Among the tested methods, Tabu Search was selected for the improvement stage due to its stability and ability to generate high-quality solutions within short computation times. It consistently produced feasible schedules with strong performance across different scenarios; see Figure 20.3 for comparison.

Validation with Real Data: The model was validated using historical weekly production data. Existing plans were reconstructed under realistic

Week	Method	Cmax	$\sum C_l$	Changeover	$\sum OT$	Max OT
W37	Original	118.43	758.14	23.60	0.00	0.00
W37	Scheduler	111.40	631.70	25.42	0.00	0.00
W37	Improve (%)	5.9	16.7	-7.7	0.0	0.0
W38	Original	168.00	966.62	28.07	144.00	48.00
W38	Scheduler	168.00	697.01	37.75	48.00	48.00
W38	Improve (%)	0.0	27.9	-34.5	66.7	0.0
W39	Original	168.00	787.18	17.25	67.44	48.00
W39	Scheduler	168.00	670.31	26.33	68.28	48.00
W39	Improve (%)	0.0	14.8	-52.6	-1.2	0.0
W40	Original	168.00	874.28	21.41	103.60	48.00
W40	Scheduler	116.68	760.79	26.58	0.00	0.00
W40	Improve (%)	30.5	13.0	-24.1	100.0	100.0
W41	Original	157.91	754.99	16.25	74.26	37.91
W41	Scheduler	128.12	692.45	24.67	19.89	8.12
W41	Improve (%)	18.9	8.3	-51.8	73.2	78.6

Figure 20.4: Comparison of original and heuristic-generated schedules

conditions to enable a fair comparison, and the detailed week-by-week validation results are presented in Figure 20.4. The comparison results show consistent improvements in the sum of line completion times, along with significant reductions in makespan and overtime in high-load weeks. In some cases, overtime was completely eliminated. While changeover times increased slightly, this reflects a controlled trade-off for improved overall efficiency.

Overall, the approach produces operationally meaningful schedules, demonstrating strong alignment with real production requirements and providing effective decision support.

20.4 Implementation and Pilot Study

The development of the Decision Support System integrates multiple software platforms to provide a practical and user-friendly planning tool. Python serves as the main backend engine, where input data is processed and the heuristics are executed. Microsoft Excel, supported by VBA, functions as the main user interface and control panel. This structure allows planners to enter weekly demand, raw material availability, and shift availabilities, and to review the generated schedules in a familiar environment. In this way, the system combines improvement capability with usability.

The system was evaluated under real operating conditions through par-

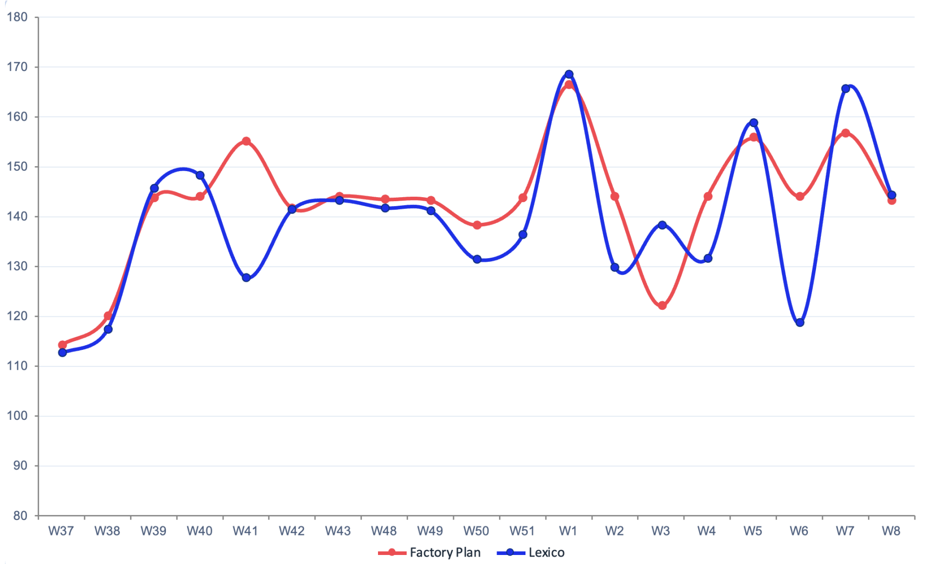

Figure 20.5: 19-week C_{max} performance comparison (12/19 wins).

allel use alongside the factory's existing manual planning process, following Unilever's approval and pilot launch on March 23. Conducted in coordination with company officials, this enabled comparison of outputs, assessment of feasibility, and identification of adjustments required by real-world conditions. The evaluation thus served not only as a performance check, but also as an opportunity to improve the tool based on user feedback.

20.5 Benchmark and Benefits

To evaluate system performance before pilot implementation, the heuristic was benchmarked against historical factory plans for 19 weeks. The results demonstrate consistent improvements. Specifically, the DSS achieved a 2.42% average reduction in C_{max}, outperforming original plans in 12 out of 19 weeks (Figure 20.5). Even more significant was the impact on sum of line completion times, where the DSS resulted in a 10.26% average reduction in $\sum_{\ell \in L} C^{\ell}$, securing superior results in 18 out of 19 weeks (Figure 20.6). On average, DSS achieved a reduction of 77 hours per week in the sum of line completion times. When operational cost factors including energy consumption, maintenance, and labor expenses are considered, the financial impact of this optimization becomes quite significant. Based on the benchmark where a single 8 hour shift reduction translates to a 1,500 € gain, the system provides a weekly saving of 14,437.50 €. Over a full 52 week production cycle, the implementation of this decision support system results in a total estimated annual profit increase of 750,750 €.

The successful implementation of this project provides Unilever with

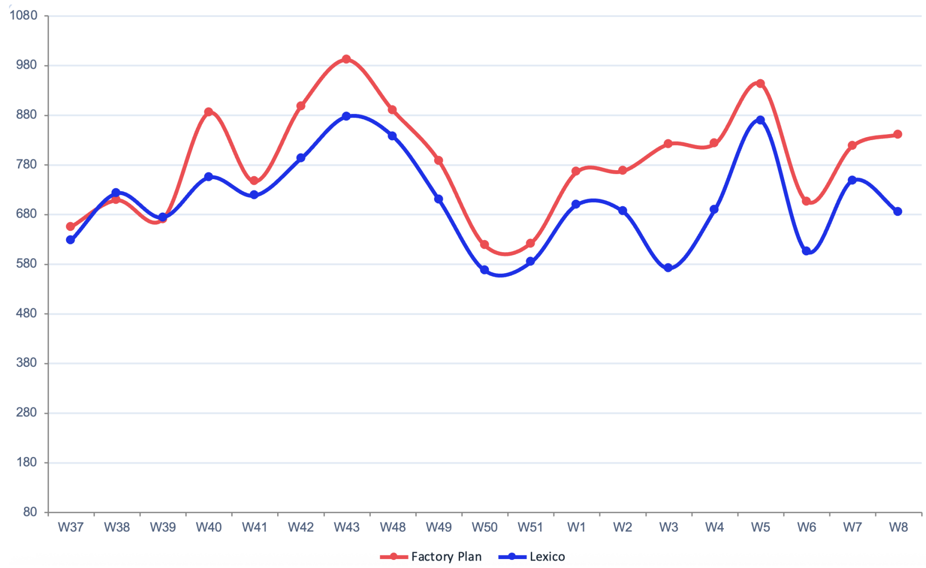

Figure 20.6: 19-week $\sum_{\ell \in L} C^{\ell}$ performance comparison(18/19 wins).

several operational and strategic benefits beyond increased profitability:

- **Bottleneck Relief:** By improving SKU sequencing and reducing the makespan, the DSS directly relieves the main bottleneck in the powder detergent packaging area. This improvement increases throughput and expands the effective capacity of the entire system.
- **Adaptive Planning Capability:** The fast execution of the DSS enables planners to rerun the system during the week, allowing for rapid adaptation to disruptions such as machine breakdowns or material delays.
- **Workforce Flexibility:** Beyond productivity gains, the significant reduction in overtime provides planners with greater flexibility in navigating legal labor regulations and weekend shift limits. This ensures that schedules remain within compliance while reducing the operational pressure on both factory staff and management.

Beyond these direct gains, transitioning from manual planning to an automated DSS establishes systematic consistency. This shift reduces the risk of human error, improving both transparency and repeatability in the planning process.

20.6 Conclusion

This project addresses a critical operational bottleneck at Unilever's Konya Home Care Factory, the largest powder detergent packaging site in Europe, by replacing expertise-based manual scheduling with a comprehensive DSS.

The framework utilizes a two-stage lexicographic mathematical model, combining a greedy heuristic with a Tabu Search improvement phase. This approach ensures that the solution is both theoretically grounded and computationally efficient for real-time environments. To validate the system's performance, the heuristic was benchmarked against 19 weeks of historical factory data. The scheduler outperformed the original manual plans in 12 of the 19 weeks for makespan ($C_{\max}$) and in 18 of the 19 weeks for total line completion times ($\sum_{\ell \in L} C^{\ell}$), achieving average reductions of 2.42% and 10.26%, respectively. These improvements translate to a projected annual profit increase of 750,750 €.

Beyond quantitative metrics, the system provides a level of structural resilience unattainable through manual planning. Its mid-week rerun capability allows planners to adapt to line failures or urgent demand fluctuations without restarting the scheduling process. The tool is delivered via a user-friendly Excel interface backed by a modular Python engine, requiring no specialized technical knowledge and allowing for future updates as factory capabilities evolve.

Given the strategic significance of the Konya site within Unilever's European network, optimizing packaging performance is vital for broader supply chain reliability. By increasing the speed and predictability of packaging operations, this system strengthens the facility's responsiveness and reinforces its position as a key production hub in Europe.

Bibliography

Glover, F. (1989). Tabu search—part i. *ORSA Journal on Computing 1*(3), 190–206.

Glover, F. (1990). Tabu search—part ii. *ORSA Journal on Computing 2*(1), 4–32.

Pinedo, M. L. (2022). *Scheduling: Theory, Algorithms, and Systems* (6th ed.). Cham: Springer.

Yapay Zeka Tabanlı Risk-Kaçıngan Üretim Çizelgelemesi | 21

SCW.AI

Proje Ekibi

Atakan Alaçam, Özgün Bilecen, Zeynep Doğan, Yılmaz Öncü, Ata Öztürk, Ataberk Sarı, Ahmet Bedirhan Soylu

Şirket Danışmanı
Haluk Atlı
Ar-Ge ve İleri Analitik Direktörü

Akademik Danışman
Prof. Dr. M. Selim Aktürk
Endüstri Mühendisliği Bölümü

Özet

Bu proje, SCW.AI'nin Dijital Fabrikası için ilişkisiz paralel makine planlama problemine veri odaklı bir çözüm önererek, belirsiz, gerçek zamanlı ortamlarda performansı iyileştirmeyi amaçlamaktadır. Mevcut sistem, geçiş sürelerindeki yüksek değişkenlik, beklenmedik makine arızaları ve belirsiz çalışma hızları gibi önemli aksaklıkları ele almakta zorlanmakta ve bu da planlama sapmalarına ve gecikmelere neden olmaktadır. Bu sorunu çözmek için, veri tabanlı sağlamcı bir planlama sistemi geliştirilmiştir. Temel çıktı, arıza sürelerini ve işlem sürelerini tahmin etmek ve çizelgeleme yapan bir sistemdir.

Anahtar Sözcükler: İlişkisiz Paralel Makine Planlaması, Sağlam Planlama, Makine Öğrenimi, Koşullu Risk Değeri

AI based Risk-Averse Production Scheduling

Abstract

This project proposes a data-driven solution to the unrelated parallel machine scheduling problem for SCW.AI's Digital Factory, aiming to improve performance in uncertain, real-time environments. The current system struggles to handle significant disruptions, such as high variability in changeover times, unexpected machine failures, and uncertain run speeds, which cause schedule deviations and tardiness. To address this, we developed a data-based robust scheduling system. The key deliverable is a standalone Decision Support System that replaces the current system for estimating downtimes and process times.

Keywords: Unrelated Parallel Machine Scheduling, Robust Scheduling, Machine Learning, Conditional Value-at-Risk.

21.1 Company Description

Established in 2014 by Evren Özkaya, Ph.D., SCW.AI began as a consulting firm focused on supply chain strategies for regulated manufacturing industries, particularly pharmaceuticals and food & beverage. To accelerate digital transformation for manufacturers, the company launched its "Digital Factory" platform in 2017. Experiencing rapid growth, SCW.AI became an independent entity in 2022, specializing in delivering this platform via a Software-as-a-Service (SaaS) model. Today, SCW.AI integrates advanced artificial intelligence and analytics to help producers automate and optimize complex decision-making tasks (SCW.AI, 2016b,a). Recognized globally as a Gartner Cool Vendor in 2019 and operating as a certified R&D center, SCW.AI has completed over 300 projects for more than 80 companies across 40 countries, driving technological advancement in the global production sector.

21.2 System Analysis and Problem

SCW.AI currently utilizes the SCW.AI Scheduler, advanced cloud-based platform primarily for pharmaceutical production planning. It features a Gantt-chart-based interface that integrates live operational data to visualize work orders, resource assignments, and machine statuses as in Figure 1. Despite its built-in rule-based algorithms, the generated schedules frequently deviate from actual production realities. Discrepancy stems from the system's reliance on deterministic master data—such as fixed catalog speeds and static setup times—which fails to account for the inherent stochasticity of the shop floor. Four main sources of variability disrupt the schedules:

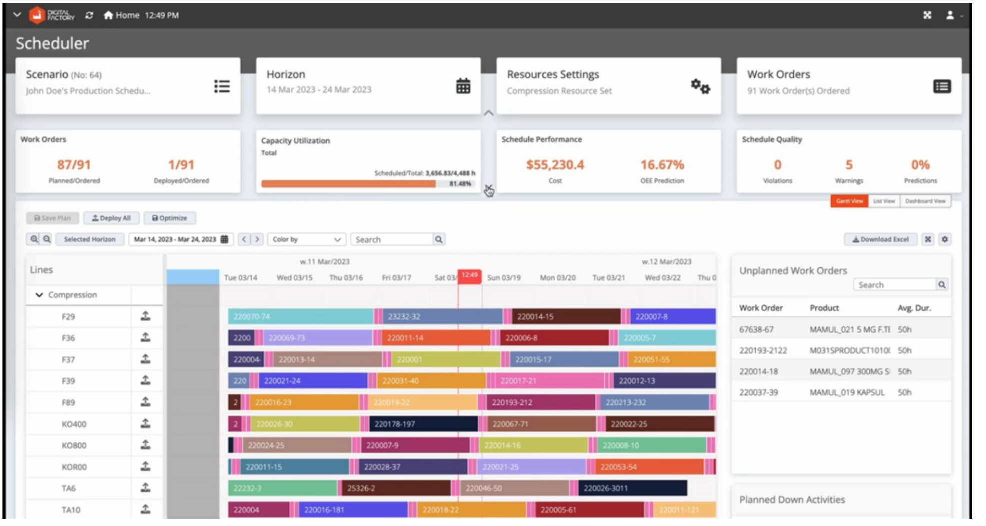

Figure 21.1: Screenshot of SCW.AI Scheduler user interface

- **Variable Transition Times:** Cleanup and setup durations fluctuate based on shift changes and operator behavior.
- **Sequence-Dependent Setups:** Changeover times heavily depend on the specific sequence of product families being processed.
- **Unexpected Machine Breakdowns:** Unplanned failures shift the entire planning horizon, and the current system lacks the probabilistic modeling to mitigate them proactively.
- **Inconsistent Run Speeds:** Machine speeds act as random variables due to mechanical aging and environmental factors, rarely sustaining ideal catalog speeds.

The combined effect of these factors creates systemic delays, reducing schedule reliability. The fundamental problem is the absence of a robust, disruption-aware scheduling mechanism capable of adapting to real-time uncertainties and generating realistic, resilient production plans.

21.3 Proposed Solution Strategy

The proposed solution strategy aims to enhance overall system efficiency and formally manage the stochasticity of the production environment. The primary objective is to minimize the expected makespan while strictly penalizing risk, maintaining a risk-averse perspective that aligns with the realities of pharmaceutical manufacturing. Ultimately, this approach provides a robust, risk-managed schedule capable of serving customers before their deadlines despite real-time disruptions.

21.3.1 Critical Assumptions

The proposed scheduling model is built upon several foundational operational assumptions. First, work orders are considered indivisible, non-preemptive tasks that require dedicated, uninterrupted setup and cleanup cycles. Second, each production line can process only one order at a time. Third, the scheduling framework operates on a continuous timeline, intentionally smoothing over standard breaks or shift changes for continuous mathematical modeling. Fourth, inventory holding costs for early completion are excluded from the objective function. Fifth, setup and cleanup durations are strictly sequence-dependent. Finally, processing and transition times are modeled using normal statistical distributions, while expected machine failures are forecasted leveraging historical data.

21.3.2 Solution Approach

Data Analysis and Parameter Estimation

To accurately mathematically model shop-floor realities, we transformed raw activity logs into structured parameters. We analyzed extensive historical data to determine four critical inputs: production run times, sequence-dependent setup and cleanup times, machine breakdown probabilities, and downtimes.

To account for real-world fluctuations, we categorized the raw data into uptimes, downtimes, setups, and cleanups. We then modeled work order run times as normal distributions , calculating a specific mean and standard deviation for each unique machine-product combination. Sequence-dependent transition times were similarly modeled. To address the inevitable gaps in historical data, we utilized a transition matrix hierarchy: if an exact product-to-product transition time was unavailable, we estimated the duration based on broader product family averages, or as a last resort, general system averages.

Machine Breakdown Estimation

To eliminate the flawed, deterministic assumption of constant machine availability, we developed predictive machine learning models to estimate both the probability and the severity of machine breakdowns. We engineered features capturing "machine fatigue," such as the exact time elapsed since the last breakdown and rolling cumulative uptime metrics.

A Logistic Regression model was trained to predict the probability of a breakdown occurring during a specific job, while a secondary Linear Regression model estimates the expected downtime duration if that failure actually occurs.

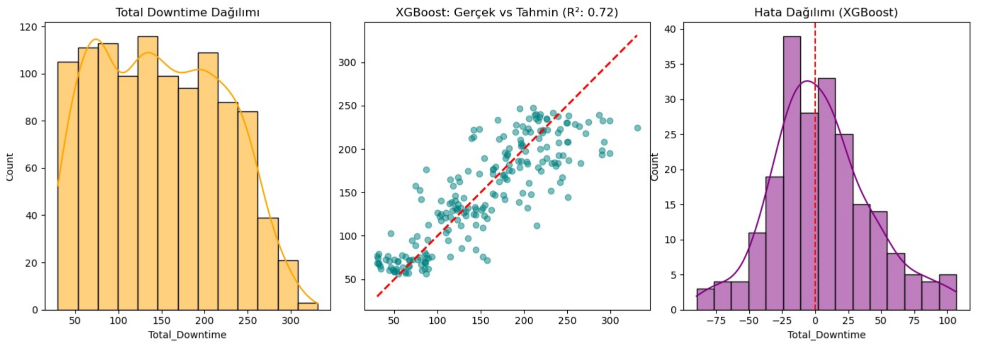

Figure 21.2: AI Model Performance Measures

Crucially, these predictions are directly integrated into the scheduling heuristic. By inflating the expected processing times for high-risk job assignments, the scheduler is forced to proactively assign larger time buffers to fatigued machines, creating a robust schedule capable of absorbing unexpected failures.

Stochastic Mathematical Model

To rigorously handle these uncertainties, we expanded our initial framework into a Stochastic Mixed-Integer Linear Program (MILP) . The model ensures every work order is assigned to exactly one production line, maintains a continuous sequence of jobs, and accurately tracks the total completion time across all active machines.

The most critical addition is the Conditional Value-at-Risk (CVaR) component. While standard deterministic optimization might suggest a schedule taking 10 hours on average, it ignores rare but severe disruptions that could push it to 20 hours. CVaR calculates the average completion time of the absolute worst percentage of possible outcomes. By mathematically penalizing this tail risk in the objective function, the scheduler actively avoids job sequences that are vulnerable to extreme delays (see Appendix A for full mathematical formulation).

Heuristic Approach

The problem involving unrelated parallel machines scheduling with sequence-dependent setup times is highly complex (NP-Hard) and practically impossible to solve exactly for large-scale manufacturing environments within a reasonable timeframe. Along with the stochastic breakdowns and processing times, the computational burden becomes overwhelming.

Therefore, we utilize a meta-heuristic: Simulated Annealing (SA). This approach intelligently searches through millions of possible schedules to find a highly robust solution quickly. It uses a specific "neighborhood structure" with three primary moves: swapping two jobs between different machines, swapping the sequence of two jobs on the same machine, or entirely reassigning a job.

- **Initialization:** The process begins with a baseline schedule derived from mean job and changeover times. To ensure a high-quality start, we utilize **Google OR-Tools**, a **gold-medal winner** in international constraint programming competitions since 2013.

- **Scenario Generation:** To test true robustness, various realistic scenarios are generated, reflecting the unpredictable nature of machine speeds, transition delays, and breakdowns modeled in previous steps.

- **Iterative Search:** The algorithm evaluates neighboring schedules across these uncertain scenarios. If a new schedule reduces expected completion time and risk, it is accepted immediately. To avoid getting trapped in local optima, worse schedules may still be accepted based on a probability tied to the current temperature.

- **Cooling Schedule:** The system's temperature is systematically reduced. As it "cools," the algorithm becomes increasingly restrictive regarding worse schedules, eventually locking into the most robust sequence discovered.

A profound component of our stochastic evaluation step is the deep integration of Clark's Approximation. Typically, evaluating makespan under uncertainty requires thousands of slow Monte Carlo simulations. However, Clark's Approximation provides a powerful analytical method to estimate the mean and variance of the maximum of several normally distributed random variables. By leveraging this technique, we are able to accomplish in milliseconds through analytical calculations what would normally require 1,000 simulations for, achieving high-fidelity results without any loss in performance.

Simulation Module

Process Times and Breakdown Generation: Because process times are statistical distributions, the output is not a static list of fixed KPIs. Instead, the module dynamically generates operational durations and disruptions. Job processing and changeover times are randomly generated using NumPy's normal distribution functions. Simultaneously, breakdowns are randomly introduced via the outcomes of our regression models, exposing the true combined impact of downtimes.
Visualisation: To help clients track stochastic schedules, a Gantt chart is

provided that integrates run times, changeovers, and simulated breakdowns simultaneously across different generated scenarios, allowing users to easily view baseline schedules and isolate late jobs.

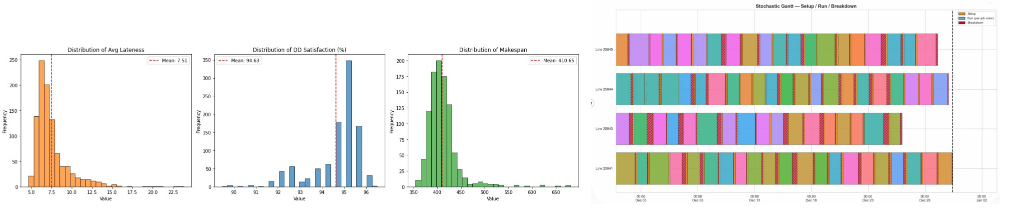

Figure 21.3: KPI Histograms and Animated Gantt Chart

21.4 Validation

The validation process was designed to rigorously test whether the proposed risk-averse scheduling algorithm practically outperforms the current operations under real-world uncertainty. Because directly testing experimental schedules on a live pharmaceutical shop floor is not currently possible, we conducted our validation using a high-fidelity stochastic simulation environment built exclusively upon historical factory data.

21.4.1 Validation of Simulation

We tested how accurate the simulation is by comparing the simulated versions of old schedules that were ran on the shop floor and the actual realization of those schedules. We observed that the introduction of stochasticity indeed decreases the difference between the actual result and the simulated schedule. We observed that the actual makespan is always within the 95% confidence interval of the generated simulated makespan.

Table 21.1: Validation for Simulation Results

Metric	**SCW Scheduler**	**SCW Scheduler Simulation**	**Actual**
December Test	689.28	925.83 ± 107.10	911.05
February Test	696.51	756.49 ± 67.82	792.98
March Test	704.27	901.57 ± 80.41	932.74

21.4.2 Measurement and Comparison Methodology

To establish a complete comparison, we evaluated the schedules based on three primary Key Performance Indicators: Total Makespan, Average Changeover per Line, and Cumulative Downtime. We measured and compared four distinct scheduling states:

- **SCW Scheduler (Theoretical):** The raw, deterministic output of the company's current planning system, assuming perfect, uninterrupted conditions.
- **SCW Scheduler Simulation (Expected Reality):** The current system's schedule evaluated through our stochastic simulation module to reveal how it actually performs under stress.
- **Actual Historical Realization:** The exact metrics recorded on the physical shop floor for the corresponding production period.
- **Our Algorithm:** The schedule generated by our Simulated Annealing heuristic, evaluated through the exact same simulation module.

By comparing the simulated performance of the current SCW Scheduler against the Actual Historical Realization, we successfully validated our simulator; the simulated metrics closely tracked actual historical records, proving our environment is a highly accurate representation of true stochastic behavior.

21.4.3 Results Analysis

The current SCW Scheduler projects an overly optimistic makespan that collapses and extends significantly when subjected to simulated disruptions. In contrast, our proposed algorithm successfully navigates these real-world variations. By intelligently inserting time buffers and avoiding sequence assignments on machines with high breakdown probabilities, our model achieves a highly executable and robust schedule. It delivers a substantial reduction in total production time compared to both the simulated baseline of the current system and actual historical performance, optimizing setup efficiency.

21.5 Integration and Implementation

21.5.1 Scheduler Platform

The Scheduler Platform provides a comprehensive, interactive graphical user interface designed specifically for production planners. The primary visualization component is a detailed Gantt chart, tracking production line allocations (vertical axis) against production time in minutes (horizontal axis). The interface is highly flexible, supporting multiple input methods: users can upload raw operational data via Excel to generate new schedules, or load pre-existing JSON and .txt files to reenact predetermined plans.

To manage interactions, a top-left toggle switches between two primary states: "Observe" mode for schedule review, and "Edit" mode for active schedule manipulation. On the right-hand panel, the platform offers dynamic control over operational parameters. Planners can adjust objective function settings based on risk tolerance: optimizing strictly for Expected Makespan, strictly for tail-risk (CVaR), or a weighted combination using a custom λ parameter. Clicking "Save Parameters" immediately reruns the mathematical model, dynamically updating metrics like machine utilization rates.

A distinctive feature is the "Animated Gantt" visualization. By utilizing playback controls, users observe the stochastic nature of manufacturing in real-time. The system continuously regenerates schedule instances, drawing job times from their statistical distributions. Planners can pause the animation to analyze specific realizations or document variations by selecting "Generate GIF."

Scheduler Playground

To deeply facilitate human-in-the-loop planning, "Edit" mode acts as an interactive Scheduler Playground. Planners are empowered to manually override the algorithm by selecting any work order and dragging it into an alternative, feasible position on the Gantt chart. The moment a manual sequence change occurs, the underlying algorithm instantly rebuilds the visual charts and triggers a fresh simulation run. By empowering planners to manually prioritize high-value orders for earlier production, our platform instantly recalculates every KPI and makespan metric, transforming subjective intuition into data-validated decision power.

21.6 Benefits to the Company

Our work provides a simulation module that can be used when deciding the optimal scheduling approaches over a given past production data. The company will be able to compare the actual makespan of the schedule which is usually extended beyond what is deterministically computed due to unexpected breakdowns and variable process and setup times. Currently, the deterministic SCW Scheduler generates a stochasticity unaware makespan of 689.28 hours. However, when subjected to real-world shop-floor realities in our simulation module, the expected makespan balloons to 925.83 hours, closely mirroring the actual historical realization of 911.05 hours which affirms the claim that the simulation module is valid.

Our algorithm replaces this fragile optimism with robust reliability. By sequencing the jobs so that the realized schedule's makespan decreases, our model achieves a highly executable expected makespan of 714.45 hours. This represents a massive reduction of over 200 hours compared to the simulated (and realized) performance of the current SCW Scheduler. The robust schedule increases the changeover times slightly, however the intentional decreasing of the possible machine breakdowns more than make up for the 8 hours of increase in changeover time and the total time is cut down by more than 200 hours.

Furthermore, because the algorithm utilizes Clark's Approximation, it evaluates complex tail-risks in under a minute, allowing SCW.AI to offer real-time rescheduling without incurring massive server compute costs. Coupled with the interactive Scheduler Platform, SCW.AI provides a cutting-edge tool that builds trust by transforming unpredictable manufacturing environments into manageable, highly optimized operations.

Table 21.2: Pilot Study Results (Makespan)

Metric	**Actual**	**SCW Scheduler Simulation**	**Our Algorithm**	**Percentage**
Pilot Study 1	911.05	925.83 ± 105.10	714.45 ± 58.64	23%
Pilot Study 2	792.98	756.49 ± 67.82	696 ± 37.65	8%
Pilot Study 3	932.74	901.57 ± 80.41	783 ± 42.65	12%

21.7 Conclusion

This project upgrades SCW.AI's scheduling system to handle the stochastic shop-floor realities. By combining machine learning with advanced optimization algorithms, our risk-averse model manages shop-floor uncertainties like sudden machine breakdowns and fluctuating transition times.

To make this complex approach practical for daily operations, the algorithms are seamlessly integrated into an interactive Scheduler Platform that allows production planners to easily visualize schedules, simulate disruptions, and manually adjust job sequences on the fly.

Rigorous benchmarking shows that this approach generates highly efficient, executable production plans that significantly outperform the previous deterministic system. Ultimately, this framework equips SCW.AI with a powerful, fast, and reliable enhancement for their Digital Factory, ensuring manufacturers can confidently meet production deadlines despite everyday operational challenges.

Bibliography

SCW.AI (2016a). About SCW.AI. https://scw.ai/company/about-us/. Accessed: 2025-10-20.

SCW.AI (2016b). AI-Driven Scheduling Solution for Greater Resilience and Profitability. https://scw.ai/product/scheduler/. Accessed: 2025-10-20.

Appendix: Stochastic Model

Sets and Parameters

$$
\begin{aligned}
J &= \text{Set of all jobs,}\\
J0 &= \text{Set of all nodes (jobs + dummy job 0)} J0 = J \cup \{0\}\\
K &= \text{Set of machines}
\end{aligned}
$$

Variables

$$
x_{jlk} = \begin{cases} 1 & \text{if there is a transition from job } j \text{ to job } l \text{ at machine } k.\\ 0 & \text{o.w.} \end{cases}
$$

u_{jk} : Order of job j at machine k

Additional sets, parameters and risk parameters

$$
\begin{aligned}
\mathcal{S} &= \{1, \ldots, S\} \quad \text{(set of scenarios)}\\
t^s_{jl} &= \text{setup time from } i \text{ to } j \text{ in scenario } s\\
p^s_{jk} &= \text{processing time of job } j \text{ on machine } k \text{ in scenario } s\\
\beta &\in (0,1) \quad \text{(CVaR confidence level)}
\end{aligned}
$$

Scenario-independent variables:

$$
\begin{aligned}
&x_{jlk} \in \{0,1\} && \forall j \in J0,\ l \in J \setminus \{j\},\ k \in K\\
&u_{jk} \in \mathbb{Z}, \quad 1 \le u_{jk} \le |J| && \forall j \in J0,\ k \in K
\end{aligned}
$$

Scenario-dependent variables:

$$
\begin{aligned}
&T^s \ge 0 && \forall s \in \mathcal{S} && \text{(makespan in scenario } s\text{)}\\
&\xi_s \ge 0 && \forall s \in \mathcal{S} && \text{(CVaR excess variable)}\\
&h \in \mathbb{R} && && \text{(CVaR linearization)}
\end{aligned}
$$

Stochastic MILP with CVaR

$$
\min \quad \frac{1}{S}\sum_{s=1}^{S} T^s \;+\; \lambda\left(h + \frac{1}{(1-\beta)S}\sum_{s=1}^{S}\xi_s\right) \tag{21.1}
$$

$$
\text{s.t.} \sum_{k\in K}\sum_{l\in J0\setminus\{j\}} x_{jlk} = 1 \qquad \forall j \in J \tag{21.2}
$$

$$\sum_{l \in J0 \setminus \{j\}} x_{ljk} - \sum_{j \in J0 \setminus \{i\}} x_{jlk} = 0 \qquad \forall j \in J0,\ k \in K \tag{21.3}$$

$$u_{jk} - u_{lk} + |J| x_{jlk} \leq |J| - 1 \qquad \forall l, j \in J,\ l \neq j,\ k \in K \tag{21.4}$$

$$\sum_{j \in J0} \sum_{l \in J \setminus \{j\}} \left(t^s_{jl} + p^s_{jk}\right) x_{jlk} \leq T^s \qquad \forall k \in K,\ s \in \mathcal{S} \tag{21.5}$$

$$u_{0k} = 1 \qquad \forall k \in I \tag{21.6}$$

$$\xi_s \geq T^s - h \qquad \forall s \in \mathcal{S} \tag{21.7}$$

$$\xi_s \geq 0 \qquad \forall s \in \mathcal{S} \tag{21.8}$$

$$x_{jlk} \in \{0, 1\}, \quad 1 \leq u_{jk} \leq |J| \qquad \forall l, j, k, s \tag{21.9}$$

$$T^s \geq 0,\ \xi_s \geq 0 \qquad \forall s \tag{21.10}$$

- **Assignment Constraint** (21.2): Each work order is performed by a production line in the factory.
- **Flow Balance Constraint** (21.3): it ensures the network solution given to the problem has correct flow-balance values depending on their respective demand and supply values.
- **MTZ Cycle Elimination Constraint** (21.4): it eliminates any cycles from forming in the network by hindering any triangles.
- **Makespan Constraint** (21.5): it ensures that the total makespan is always bounded below by all (consequently the maximum) of the individual time elapsed for production in each line.
- **Initialization Constraint** (21.6): it ensures that the model starts ordering from index 0.
- **CVaR Linearization Constraints** (21.7)-(21.1): they linearize the Value at Risk constraints. The x_{jlk}, u_{jk} decisions are *here-and-now* (common for all scenarios), while T^s, ξ_s and the CVaR threshold t are scenario-linked / auxiliary.
- **Domain Constraint** (21.9) - (21.10): they define the domain constraints.

Dizin

Index